Journal of Machine Learning Research 1 (2000) 1-48 Submitted 4/00; Published 10/00

The Mechanism of Additive Composition

Ran Tian TIANRAN@ECEI.TOHOKU.AC.JP
Naoaki Okazaki OKAZAKI@ECEI.TOHOKU.AC.JP
Kentaro Inui INUI@ECEI.TOHOKU.AC.JP
Communication Science Laboratory
Tohoku University
6-6-05 Sendai, Miyagi 980-8579, Japan

Editor:

arXiv:1511.08407v3 [cs.CL] 6 Jun 2016

Abstract

Additive composition (Foltz et al., 1998; Landauer and Dutnais, 1997; Mitchell and Lapata, 2010) is a widely used method for computing meanings of phrases, which takes the average of vector representations of the constituent words. In this article, we prove an upper bound for the bias of additive composition, which is the first theoretical analysis on compositional frameworks from a machine learning point of view. Our proof relies on properties of natural language data that are empirically verified, and can be theoretically derived from an assumption that the data is generated from a Hierarchical Pitman-Yor Process. The theory endorses additive composition as a reasonable operation for calculating meanings of phrases, and suggests ways to improve additive compositionality, including: transforming entries of distributional word vectors by a function that meets a specific condition, constructing a novel type of vector representations to make additive composition sensitive to word order, and utilizing singular value decomposition to train word vectors.

Keywords: Compositional Distributional Semantics, bias and variance, approximation error bounds, natural language data, Hierarchical Pitman-Yor Process

1. Introduction

The decomposition of generalization errors into bias and variance (Geman et al., 1992) is one of the most profound insights of learning theory. Bias is caused by low capacity of models when the training samples are assumed to be infinite, whereas variance is caused by overfitting to finite samples. In this article, we apply the analysis to a new set of problems in Compositional Distributional Semantics, which studies the calculation of meanings of natural language phrases by vector representations of their constituent words. We prove an upper bound for the bias of a widely used compositional framework, the additive composition (Foltz et al., 1998; Landauer and Dutnais, 1997; Mitchell and Lapata, 2010).

Calculations of meanings are fundamental problems in Natural Language Processing (NLP). In recent years, vector representations have seen great success at conveying meanings of individual words (Levy et al., 2015). These vectors are constructed from statistics of contexts surrounding the words, based on the Distributional Hypothesis that words occurring in similar contexts tend to have similar meanings (Harris, 1954). For example, given a target word t, one can consider its context as close neighbors of t in a corpus, and assess the probability p_i^t of the i-th word (in a fixed lexicon) occurring in the context of t. Then,

©2000 Ran Tian, Naoaki Okazaki and Kentaro Inui.

the word t is represented by a vector $\left(F(p_i^t)\right)_i$ (where F is some function), and words with similar meanings to t will have similar vectors (Miller and Charles, 1991).

Beyond the word level, a naturally following challenge is to represent meanings of phrases or even sentences. Based on the Distributional Hypothesis, it is generally believed that vectors should be constructed from surrounding contexts, at least for phrases observed in a corpus (Boleda et al., 2013). However, a main obstacle here is that phrases are far more sparse than individual words. For example, in the British National Corpus (BNC) (The BNC Consortium, 2007), which consists of 100M word tokens, a total of 16K lemmatized words are observed more than 200 times, but there are only 46K such bigrams, far less than the 16000^2 possibilities for two-word combinations. Be it a larger corpus, one might only observe more rare words due to Zipf's Law, so most of the two-word combinations will always be rare or unseen. Therefore, a direct estimation of the surrounding contexts of a phrase can have large sampling error. This partially fuels the motivation to construct phrase vectors from combining word vectors (Mitchell and Lapata, 2010), which also bases on the linguistic intuition that meanings of phrases are "composed" from the meanings of their constituent words. From a machine learning point of view, word vectors have smaller sampling error, or lower variance, because words are more abundant than phrases. Hence, a compositional framework that calculates meanings from word vectors will be favorable, if the bias of this composition operation is also small.

Here, "bias" is the distance between two types of phrase vectors, one calculated from composing the vectors of constituent words (composed vector), and the other assessed from context statistics where the phrase is treated as a target (natural vector). The statistics is assessed from an infinitely large *ideal corpus*, so that the natural vector of the phrase can be reliably estimated without sampling error, hence conveying the meaning of the phrase by Distributional Hypothesis. If the distance between the two vectors is small, the composed vector can be viewed as a reasonable approximation of the natural vector, hence an approximation of meaning; moreover the composed vector can be more reliably estimated from finite *real corpora* because words are more abundant than phrases. Therefore, an upper bound for the bias will provide a learning-theoretic support for the composition operation.

A number of compositional frameworks have been proposed in the literature (Baroni and Zamparelli, 2010; Grefenstette and Sadrzadeh, 2011; Socher et al., 2012; Paperno et al., 2014; Hashimoto et al., 2014). Some are complicated methods based on linguistic intuitions (Coecke et al., 2010), and others are compared to human judgments for evaluation (Mitchell and Lapata, 2010). However, none of them has been previously analyzed regarding their bias[1]. The most widely used framework is the additive composition (Foltz et al., 1998; Landauer and Dutnais, 1997), in which the composed vector is calculated by averaging word vectors. Yet, it was unknown if this average is by any means related to statistics of contexts surrounding the corresponding phrases.

In this article, we prove an upper bound for the bias of additive composition of two-word phrases, and demonstrate several applications of the theory. An overview is given as below.

1. Unlike natural vectors which always lie in the same space as word vectors, some compositional frameworks construct meanings of phrases in different spaces. Nevertheless, we argue that even in such cases it is reasonable to require some mappings to a common space, because humans can usually compare meanings of a word and a phrase. Then, by considering distances between mapped images of composed vectors and natural vectors, we can define bias and call for theoretical analysis.

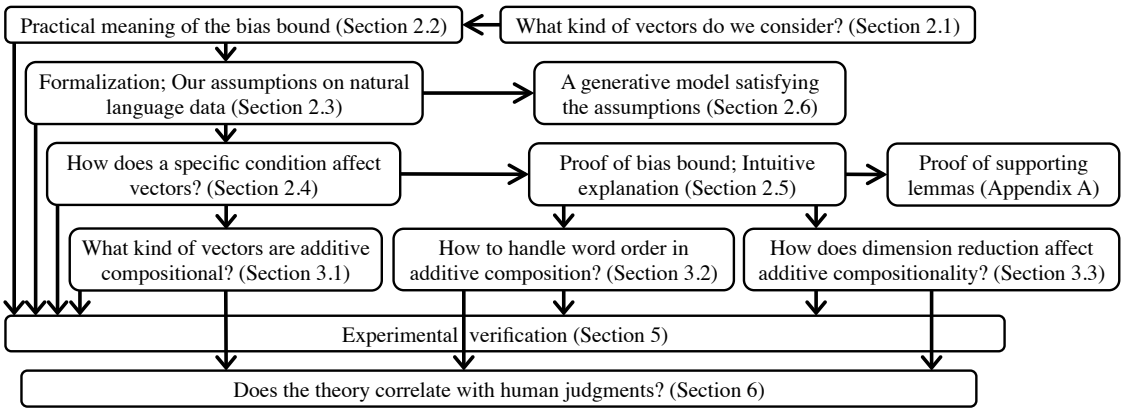
Practical meaning of the bias bound (Section 2.2)
What kind of vectors do we consider? (Section 2.1)
Formalization; Our assumptions on natural language data (Section 2.3)
A generative model satisfying the assumptions (Section 2.6)
How does a specific condition affect vectors? (Section 2.4)
Proof of bias bound; Intuitive explanation (Section 2.5)
Proof of supporting lemmas (Appendix A)
What kind of vectors are additive compositional? (Section 3.1)
How to handle word order in additive composition? (Section 3.2)
How does dimension reduction affect additive compositionality? (Section 3.3)
Experimental verification (Section 5)
Does the theory correlate with human judgments? (Section 6)

Figure 1: An overview of this article. Arrows show dependencies between sections.

In Section 2.1, we introduce notations and define the vectors we consider in this work.

In Section 2.2, we describe our bias bound for additive composition, focusing on its practical consequences that can be tested on a natural language corpus. The consequences are experimentally verified in Section 5.2.

In Section 2.3, we give a formalized version of the bias bound (Theorem 1), with our assumptions on natural language data clarified. These assumptions include the well-know Zipf's Law, a generalized version of the law applied to word co-occurrences, and some intuitively acceptable conditions. The assumptions and their implications are experimentally verified in Section 5. Moreover, we show that the Generalized Zipf's Law holds if the data is generated by a Hierarchical Pitman-Yor Process (Section 2.6).

One of our theoretical findings is that the function F applied to entries of vector representations is required to meet a specific condition. In Section 2.4, we explain how this condition affects vector representations and why it is important to the bias bound.

In Section 2.5, we sketch the proof of our bias bound and give an intuitive explanation. Proofs of some supporting lemmas are left to Appendix A.

We demonstrate three applications of our theory:

1. The condition required to be satisfied by F provides a unified explanation on why some recently proposed word vectors are good at additive composition (Section 3.1). Our experiments also verify that the condition drastically affects additive compositionality and other properties of vector representations (Section 5.2, Section 6).

2. Our intuitive explanation for the bias bound inspires a novel method for making vectors recognize word order, which was long thought as an issue for additive composition (Section 3.2, Section 5.3). Unlike previous approaches, we still compose vectors by taking average, retaining the merits of being parameter-free and having a bias bound.

3. Our theory suggests that Singular Value Decomposition (SVD) is suitable for preserving additive compositionality in dimension reduction of word vectors (Section 3.3). Experiments also show that SVD might perform better than other models in additive composition (Section 5.4, Section 6).

A summary is illustrated in Figure 1.

*... as a percentage of your income, your **tax rate** is generally less than that ...*

target	words in context
tax	*percentage, of, your, income, your, rate, is, generally, less, than*
rate	*of, your, income, your, tax, is, generally, less, than, that*
tax_rate	*of, your, income, your, is, generally, less, than*

Table 1: Contexts are taken as the closest five words to each side for the targets "*tax*" and "*rate*", and four for the target "*tax_rate*".

2. Theory

In this section, we discuss vector representations constructed from an ideal natural language corpus, and establish a mathematical framework for analyzing additive composition.

2.1 Notation and Vector Representation

A natural language corpus is a sequence of words. Ideally, we assume that the sequence is infinitely long and contains an infinite number of distinct words.

Notation 1. In practice, only finite restrictions of the ideal corpus are observed. We denote the observed number of distinct words by n, and use the n words as a lexicon to construct vector representations. Then, we will consider the limit $n \to \infty$. Let $p_{i,n}$ be the observed probability of the i-th word ($1 \leq i \leq n$). Assume i is taken such that $p_{i,n} \geq p_{i+1,n}$.

With a corpus given, we consider vector representations for *targets*, which are either words or phrases. To define the vectors we start from specifying a *context* for each target, which is usually taken as words surrounding the target in the corpus. Table 1 shows an example of a sequence of words, and the contexts of a phrase target and two word targets. The contexts are taken as words close to the targets in the sequence.

Notation 2. We use s, t to denote word targets, and st a phrase target consisting of two consecutive words s and t. When the word order is ignored (i.e., either st or ts), we denote the target by $\{st\}$. A general target is denoted by Υ. Later in this article, we will consider other types of targets as well, and a full list of target types is shown in Table 2.

Notation 3. Let $p^{\Upsilon}_{i,n}$ denote the co-occurrence probability of the i-th word, conditioned on it being in the context of Υ and the lexicon size being n.

Definition 4. In this work, we consider vector representations of the form

$$\mathbf{w}^{\Upsilon}_n := \big(w^{\Upsilon}_{i,n}\big)_{1\leq i\leq n}, \quad \text{where} \quad w^{\Upsilon}_{i,n} = c_n\big(F(p^{\Upsilon}_{i,n}+1/n) - a^{\Upsilon}_n - b_{i,n}\big).$$

In above, a^{Υ}_n, $b_{i,n}$ and c_n are real numbers and F is a smooth function on $(0,\infty)$. In Section 2.2, we will further specify a^{Υ}_n, $b_{i,n}$, c_n and F without much loss of generality.

To consider $F(p^{\Upsilon}_{i,n}+1/n)$ instead of $F(p^{\Upsilon}_{i,n})$ is a smoothing scheme that guarantees $F(x)$ being applied to $x > 0$. We will consider F that is not continuous at 0, such as $F(x) := \ln x$; yet, $w^{\Upsilon}_{i,n}$ has to be well-defined even if $p^{\Upsilon}_{i,n} = 0$. In practice, the $p^{\Upsilon}_{i,n}$ estimated from a finite corpus can often be 0; theoretically, the smoothing scheme plays a role in our proof as well.

Notation 2	Υ	a general target can denote either of the following:
Notation 2	s, t	word targets
Notation 2	st	two-word phrase target
Notation 2	$\{st\}$	two-word phrase target with word order ignored
Definition 6	$s/t\backslash s$	a word t not neighboring the word s
Definition 15	$s\bullet, \bullet t$	words in the context of $s\bullet$ (resp. $\bullet t$) are assigned the left-hand-side (resp. right-hand-side) Near-far labels
Definition 16	$s\bullet\backslash t, s/\bullet t$	a target $s\bullet$ (resp. $\bullet t$) not at the left (resp. right) of word t (resp. s)
Theorem 1	S, T	random word targets
General	random word targets can form different types such as $\{ST\}$ and $S/T\backslash S$	

Table 2: List of target types

The definition of $\mathbf{w}_n^\Upsilon$ is general enough to cover a wide range of previously proposed distributional word vectors. For example, if $F(x) = \ln x$, $a_n^\Upsilon = 0$ and $b_{i,n} = \ln p_{i,n}$, then $w_{i,n}^\Upsilon$ is the Point-wise Mutual Information (PMI) value that has been widely adopted in NLP (Church and Hanks, 1990; Dagan et al., 1994; Turney, 2001; Turney and Pantel, 2010). More recently, the Skip-Gram with Negative Sampling (SGNS) model (Mikolov et al., 2013b) is shown to be a matrix factorization of the PMI matrix (Levy and Goldberg, 2014b); and the more general form of a_n^Υ and $b_{i,n}$ is explicitly introduced by the GloVe model (Pennington et al., 2014). Regarding other forms of F, it has been reported in Lebret and Collobert (2014) and Stratos et al. (2015) that empirically $F(x) := \sqrt{x}$ outperforms $F(x) := x$. We will discuss this topic further in Section 3.1.

Finally, the scalar c_n is a normalization factor that controls the scales of vectors.

We finish this section by pointing to Table 3 for a list of frequently used notations.

2.2 Practical Meaning of the Bias Bound

A compositional framework combines vectors $\mathbf{w}_n^s$ and $\mathbf{w}_n^t$ to represent the meaning of phrase "$s\ t$". In this work, we study relations between this composed vector and the natural vector $\mathbf{w}_n^{\{st\}}$ of the phrase target[2]. More precisely, we study the Euclidean distance

$$\lim_{n\to\infty} \|\mathbf{w}_n^{\{st\}} - \mathrm{COMP}(\mathbf{w}_n^s, \mathbf{w}_n^t)\|$$

where $\mathrm{COMP}(\cdot,\cdot)$ is the composition operation. At limit $n \to \infty$, the corpus is infinitely large, so $\mathbf{w}_n^{\{st\}}$ can be well estimated to represent the meaning of "$s\ t$ or $t\ s$". Therefore, the above distance can be viewed as the bias of approximating $\mathbf{w}_n^{\{st\}}$ by the composed vector $\mathrm{COMP}(\mathbf{w}_n^s, \mathbf{w}_n^t)$. In practice, especially when COMP is a complicated operation with parameters, it has been a widely adopted approach to learn the parameters by minimizing the same distances for phrases observed in a corpus (Dinu et al., 2013; Baroni and Zamparelli, 2010; Guevara, 2010). These practices further motivate our study on the bias.

2. Or it should be $\mathbf{w}_n^{st}$ if one cares about word order, which we will discuss in Section 3.2.

Notation 1	i, n	index $1 \le i \le n$, where n is the lexicon size
Notation 1	$p_{i,n}$	probability of the i-th word, $p_{i,n} \ge p_{i+1,n}$
Notation 3	$p_{i,n}^{\Upsilon}$	probability of the i-th word conditioned on it being in the context of Υ and lexicon size being n
Definition 4	$\mathbf{w}_n^{\Upsilon} := \big(w_{i,n}^{\Upsilon}\big)_{1\le i\le n}$, where $w_{i,n}^{\Upsilon} = c_n\big(F(p_{i,n}^{\Upsilon}+1/n) - a_n^{\Upsilon} - b_{i,n}\big)$	
Definition 5	$\mathcal{B}_n^{\{st\}} := \|\mathbf{w}_n^{\{st\}} - \frac{1}{2}(\mathbf{w}_n^s + \mathbf{w}_n^t)\|$	
Definition 6	$\pi_{s/t\backslash s}$	probability for an occurrence of t being non-neighbor of s
Definition 7	Λ_n	set of observed two-word phrases, word order ignored
General	$\mathbb{E}[\cdot], \mathrm{Var}[\cdot]$	expected value and variance of a random variable
General	$I_{\mathcal{H}}$	indicator; $I_{\mathcal{H}} = 1$ if condition $\mathcal{H}$ is true, 0 otherwise
General	$\mathbb{P}(\mathcal{H})$	probability of $\mathcal{H}$ being true; $\mathbb{P}(\mathcal{H}) = \mathbb{E}[I_{\mathcal{H}}]$
General	$\lambda, \beta, \xi, \ldots$	lowercase Greek letters denote real constants
Theorem 1	$X := p_{i,n}^{\Upsilon}/p_{i,n}$ where $\Upsilon := \{ST\}$, $S/T\backslash S$, or $T/S\backslash T$.	
Lemma 1	$Y_{i,n} := F(p_{i,n}X+1/n) - F(p_{i,n}\beta+1/n)$	
Lemma 1	$\varphi_{i,n} := (p_{i,n})^{2\lambda}\big(1 + (\beta n p_{i,n})^{-1}\big)^{-1+2\lambda}$	

Table 3: Frequently used notations and general conventions in this article

Definition 5. Consider the additive composition, where $\mathrm{COMP}(\mathbf{w}_n^s, \mathbf{w}_n^t) := \frac{1}{2}(\mathbf{w}_n^s + \mathbf{w}_n^t)$ is a parameter-free operation. We define

$$\mathcal{B}_n^{\{st\}} := \|\mathbf{w}_n^{\{st\}} - \frac{1}{2}(\mathbf{w}_n^s + \mathbf{w}_n^t)\|.$$

Our analysis starts from the observation that, every word in the context of $\{st\}$ also occurs in the contexts of s and t: as illustrated in Table 1, if a word token t (e.g. "*rate*") comes from a phrase $\{st\}$ (e.g. "*tax rate*"), and if the context window size is not too small, the context for this token of t is almost the same as the context of $\{st\}$. This motivates us to decompose the context of t into two parts, one coming from $\{st\}$ and the other not.

Definition 6. Define target $s/t\backslash s$ which counts every occurrence of word t not neighboring s. We use $\pi_{s/t\backslash s}$ to denote the probability of t not neighboring s, conditioned on an occurrence of t. Practically, $(1 - \pi_{s/t\backslash s})$ can be estimated by the ratio $C(\{st\})/C(t)$, where $C(\cdot)$ denotes frequency count. Then, we have the following equation

$$p_{i,n}^t = \pi_{s/t\backslash s} p_{i,n}^{s/t\backslash s} + (1 - \pi_{s/t\backslash s}) p_{i,n}^{\{st\}} \quad \text{for all } i, n \tag{1}$$

because a word in the context of t occurs in the context of either $\{st\}$ or $s/t\backslash s$.

We can view $\pi_{s/t\backslash s}$ (and similarly $\pi_{t/s\backslash t}$) as indicating how weak the "collocation" between s and t is. When $\pi_{s/t\backslash s}$ and $\pi_{t/s\backslash t}$ are small, s tends to occur with t together, so we expect $\mathbf{w}_n^s$ and $\mathbf{w}_n^t$ to be correlated with $\mathbf{w}_n^{\{st\}}$, making $\mathcal{B}_n^{\{st\}}$ small. This is exactly how we bound $\mathcal{B}_n^{\{st\}}$. However, to make $\mathcal{B}_n^{\{st\}}$ comparable we need normalize a_n^{Υ}, $b_{i,n}$, c_n and F.

To normalize a_n^Υ, $b_{i,n}$, c_n and F, we first note that $\mathcal{B}_n^{\{st\}}$ does not depend on $b_{i,n}$ because it is canceled out in the following:

$$\big(\mathcal{B}_n^{\{st\}}\big)^2 = c_n^2 \sum_{i=1}^{n} \Big(\big(F(p_{i,n}^{\{st\}}+1/n) - \frac{F(p_{i,n}^{s}+1/n) + F(p_{i,n}^{t}+1/n)}{2}\big) - (a_n^{\{st\}} - \frac{a_n^s + a_n^t}{2})\Big)^2.$$

Thus, $b_{i,n}$ can be arbitrarily fixed without loss of generality. We use $b_{i,n}$ to move the centroid of natural phrase vectors to $\mathbf{0}$, so that they will not concentrate to one direction.

Definition 7. Let Λ_n be the set of two-word phrases observed in a finite corpus, word order ignored. We put

$$b_{i,n} := \frac{1}{|\Lambda_n|} \sum_{\{st\}\in\Lambda_n} F(p_{i,n}^{\{st\}}+1/n) - a_n^{\{st\}}, \tag{2}$$

so that

$$\frac{1}{|\Lambda_n|} \sum_{\{st\}\in\Lambda_n} \mathbf{w}_n^{\{st\}} = \mathbf{0}.$$

Next, note that if c_n and F are fixed, $\mathcal{B}_n^{\{st\}}$ is taking minimum at

$$a_n^{\{st\}} - \frac{a_n^s + a_n^t}{2} = \frac{1}{n}\sum_{i=1}^{n} F(p_{i,n}^{\{st\}}+1/n) - \frac{F(p_{i,n}^{s}+1/n) + F(p_{i,n}^{t}+1/n)}{2},$$

so it is favorable to have the equality. We achieve this by adjusting a_n^Υ such that the entries of each vector average to 0.

Definition 8. We set

$$a_n^\Upsilon := \frac{1}{n}\sum_{i=1}^{n} F(p_{i,n}^{\Upsilon}+1/n) - b_{i,n}, \tag{3}$$

so that

$$\frac{1}{n}\sum_{i=1}^{n} w_{i,n}^{\Upsilon} = 0 \quad \text{for all } \Upsilon.$$

One can calculate a_n^Υ and $b_{i,n}$ by first assuming $b_{i,n} = 0$ in (3) to obtain a_n^Υ, and then substitute $a_n^{\{st\}}$ in (2) to obtain the actual $b_{i,n}$. The value of a_n^Υ will not change because if all vectors have average entry 0, so dose their centroid.

Definition 9. For c_n, we note that it changes the scales of both $\mathcal{B}_n^{\{st\}}$ and $\|\mathbf{w}_n^{\{st\}}\|$. So we set c_n such that the average norm of natural phrase vectors becomes 1:

$$\frac{1}{|\Lambda_n|} \sum_{\{st\}\in\Lambda_n} \|\mathbf{w}_n^{\{st\}}\| = 1.$$

Note that, if the centroid of natural phrase vectors had been far from $\mathbf{0}$, the normalization in Definition 9 would have caused all phrase vectors cluster around one point on the unit sphere. Then, the phrase vectors would not have been able to distinguish different meanings of phrases. The choice of $b_{i,n}$ in Definition 7 prevents such degenerated cases.

In Section 2.4, we will calculate the asymptotic values of a_n^Υ, $b_{i,n}$ and c_n theoretically.

Definition 10. For simplicity, we consider F such that $F'(x) = x^{-1+\lambda}$ in this article. So $F(x)$ can be x^λ/λ (where $\lambda \neq 0$) or $\ln x$ (where $\lambda = 0$). But intuitively, only the behavior of $F(x)$ at $x \approx 0$ matters, because F is applied to probability values close to 0. In fact, our results can be generalized to $G(x)$ such that

$$\lim_{x\to 0} G'(x)x^{1-\lambda} = 1 \quad \text{and} \quad G'(x)x^{1-\lambda} \leq M \text{ for some constant } M.$$

Remark 6 in Section 2.3 gives a formal explanation. The exponent λ turns out to be a crucial factor in our theory.

Now, we can state a summarized version of our bias bound.

Claim 1. *Assume $\lambda < 0.5$, and distributional vectors are constructed from an ideal natural language corpus. Then:*

$$\lim_{n\to\infty} \mathcal{B}_n^{\{st\}} \leq \sqrt{\frac{1}{2}(\pi_{s/t\backslash s}^2 + \pi_{t/s\backslash t}^2 + \pi_{s/t\backslash s}\pi_{t/s\backslash t})}.$$

As we expected, for more "collocational" phrases, since $\pi_{s/t\backslash s}$ and $\pi_{t/s\backslash t}$ are smaller, the upper bound for $\mathcal{B}_n^{\{st\}}$ becomes stronger. Claim 1 states a prediction that can be empirically tested on a real large corpus; namely, one can estimate $p_{i,n}^\Upsilon$ from the corpus and construct $\mathbf{w}_n^\Upsilon$ for a fixed n, then check if the inequality holds approximately while omitting the limit. In Section 5.2, we conduct the experiment and verify the prediction. In Section 2.3, we will specify the theoretical assumptions we put on an "ideal natural language corpus".

Besides it being empirically verified for phrases observed in a real corpus, the true value of Claim 1 is that the upper bound holds for an arbitrarily large ideal corpus. We can assume any plausible two-word phrase to occur sufficiently many times in the ideal corpus, even when it is unseen in the real one. In that case, a natural vector for the phrase can only be reliably estimated from the ideal corpus, but Claim 1 suggests that additive composition of word vectors provides a reasonable approximation for that unseen natural vector. Meanwhile, since word vectors *can* be reliably estimated from the real corpus, Claim 1 endorses additive composition as a reasonable meaning representation for unseen or rare phrases. On the other hand, it endorses additive composition for frequent phrases as well, because such phrases usually have strong collocations and Claim 1 says that the bias in this case is small.

The following is an outstanding by-product of our theory (Theorem 2 in Section 2.4), and it is empirically verified in Section 5.2 as well.

Claim 2. *Under the same conditions as in Claim 1, we have:*

$$\lim_{n\to\infty} \|\mathbf{w}_n^{\{st\}}\| = 1 \quad \textit{for all } \{st\}.$$

So, approximately all natural phrase vectors are distributed on the unit sphere. This makes it possible to convert the Euclidean distance $\mathcal{B}_n^{\{st\}}$ into cosine similarity, which is the most widely used similarity measure in practice.

2.3 Formalization and Assumptions on Natural Language Data

Claim 1 is formalized as Theorem 1 in the following.

Theorem 1. *For an ideal natural language corpus, we assume that:*

(A) $\lim_{n\to\infty} p_{i,n} \cdot i \ln n = 1$.

(B) *Let S, T be randomly chosen word targets. Consider $\Upsilon := \{ST\}$, $S/T\backslash S$, or $T/S\backslash T$.*

(B1) *For each n, it is as if $p_{i,n}^{\Upsilon}$ $(1 \le i \le n)$ are independent random variables.*

(B2) *Put $X := p_{i,n}^{\Upsilon}/p_{i,n}$. There exist ξ, β such that $\mathbb{P}(x \le X) = \xi/x$ for $x \ge \beta$.*

(C) *For each i and n, the random variables $p_{i,n}^{S/T\backslash S}$ and $p_{i,n}^{T/S\backslash T}$ are independent, whereas $F(p_{i,n}^{S/T\backslash S}+1/n)$ and $F(p_{i,n}^{\{ST\}}+1/n)$ have positive correlation.*

Then, if $\mathbb{E}\big[F(X)^2\big] < \infty$, we have

$$\lim_{n\to\infty} \mathcal{B}_n^{\{ST\}} \le \sqrt{\frac{1}{2}(\pi_{S/T\backslash S}^2 + \pi_{T/S\backslash T}^2 + \pi_{S/T\backslash S}\pi_{T/S\backslash T})} \quad \textit{in probability.}$$

We explain the assumptions of Theorem 1 in details below.

Remark 1. Assumption (A) is the Zipf's Law (Zipf, 1935), which states that the frequency of the i-th word is inversely proportional to i. So $p_{i,n}$ is proportional to i^{-1}, and the factor $\ln n$ comes from $\sum_{i=1}^n p_{i,n} = 1$ and $\sum_{i=1}^n i^{-1} \approx \ln n$. One immediate implication of Zipf's Law is that $np_{i,n}$ can be arbitrarily small. More precisely, for any $\delta > 0$, we have

$$np_{i,n} \le \delta \Leftrightarrow \frac{n}{\delta \ln n} \le i \le n. \tag{4}$$

The limit $np_{i,n} \to 0$ will be extensively explored in our theory.

Empirically, Zipf's Law has been thoroughly tested under several settings (Montemurro, 2001; Ha et al., 2002; Clauset et al., 2009; Corral et al., 2015).

Remark 2. When a target Υ is randomly chosen, the probability value $p_{i,n}^{\Upsilon}$ shows randomness. The assumption (B1) formalizes this intuition. Further, (B2) suggests that $p_{i,n}^{\Upsilon}$ is in the same scale as $p_{i,n}$, and $X := p_{i,n}^{\Upsilon}/p_{i,n}$ has a power law tail of index 1. We regard it as the *Generalized Zipf's Law* because it is parallel to Zipf's Law in that $p_{i,n}$ also has a power law tail of index 1. In Section 2.6, we show that Assumption (B) is closely related to a Hierarchical Pitman-Yor Process; and in Section 5.1 we empirically verify this assumption.

The assumption (B2) on power law tail can be relaxed to $\lim_{x\to\infty} x\mathbb{P}(x \le X) = \xi$. We only consider (B2) for simplicity.

Remark 3. Assumption (C) is based on the observation that generally, $S/T\backslash S$ and $T/S\backslash T$ are different word targets; and by definition, $p_{i,n}^{S/T\backslash S}$ and $p_{i,n}^{T/S\backslash T}$ are assessed from disjoint parts of the corpus. Therefore, it is natural to assume that $p_{i,n}^{S/T\backslash S}$ and $p_{i,n}^{T/S\backslash T}$ are independent. On the other hand, the targets $S/T\backslash S$ and $\{ST\}$ both contain a word T, so we expect $F(p_{i,n}^{S/T\backslash S}+1/n)$ and $F(p_{i,n}^{\{ST\}}+1/n)$ to have positive correlation.

Remark 4. The existence of second moment $\mathbb{E}\big[F(X)^2\big] < \infty$ requires $\lambda < 0.5$, because X has a power law tail of index 1 by Assumption (B2). Conversely, $\lambda < 0.5$ is usually a sufficient condition for $\mathbb{E}\big[F(X)^2\big] < \infty$, for instance, if X follows the Pareto Distribution (i.e. $\xi = \beta$) or Inverse-Gamma Distribution. Another example is given in Section 2.6.

Lemma 1. *Consider* $Y_{i,n} := F(p_{i,n}^{\Upsilon}+1/n) - F(p_{i,n}\beta+1/n) = F(p_{i,n}X+1/n) - F(p_{i,n}\beta+1/n)$. *Put* $e_{i,n} := \mathbb{E}\big[Y_{i,n}\big]$, $v_{i,n} := \mathrm{Var}\big[Y_{i,n}\big]$, *and* $\varphi_{i,n} := (p_{i,n})^{2\lambda}\big(1 + (\beta n p_{i,n})^{-1}\big)^{-1+2\lambda}$. *Then,*

(a) *There exists* χ *such that* $\dfrac{|e_{i,n}|}{\sqrt{\varphi_{i,n}}} \le \chi$ *for all* i, n.

(b) $\lim\limits_{np_{i,n}\to 0} \dfrac{e_{i,n}}{\sqrt{\varphi_{i,n}}} = 0$.

(c) *The set of random variables* $\big\{Y_{i,n}^2/\varphi_{i,n}\big\}$ *is uniformly integrable; i.e., for any* $\varepsilon > 0$, *there exists* N *such that* $\mathbb{E}\big[Y_{i,n}^2 I_{Y_{i,n}^2 > N\varphi_{i,n}}\big] < \varepsilon\varphi_{i,n}$ *for all* i, n.

(d) $\lim\limits_{np_{i,n}\to 0} \dfrac{v_{i,n}}{\varphi_{i,n}} = \eta \neq 0$, *where* $\eta = \displaystyle\int_0^\infty \big(F(z+\beta) - F(\beta)\big)^2 \cdot \frac{\xi\,\mathrm{d}z}{z^2}$.

Remark 5. Lemma 1 is the basis of our analysis. It shows that, as i and n vary, $e_{i,n}^2$ and $v_{i,n}$ scale with $\varphi_{i,n}$. Note that for random variable X we only assume a power law tail; the full distribution of X might depend on i, n and whether $\Upsilon = \{ST\}$, $S/T\backslash S$, or $T/S\backslash T$. Nevertheless, Lemma 1(b)(d) suggests that at limit $np_{i,n} \to 0$, the behavior of $e_{i,n}/\sqrt{\varphi_{i,n}}$ and $v_{i,n}/\varphi_{i,n}$ do not depend on i, n or Υ; the power law tail of X dominates.

Remark 6. The function $F(x)$ in Lemma 1 can be generalized to function $G(x)$ as mentioned in Definition 10. Because by Cauchy's Mean Value Theorem,

$$\frac{G(p_{i,n}x+1/n) - G(p_{i,n}\beta+1/n)}{F(p_{i,n}x+1/n) - F(p_{i,n}\beta+1/n)} = G'(\zeta)\zeta^{1-\lambda} \quad \text{for some } p_{i,n}\beta+1/n \le \zeta \le p_{i,n}x+1/n,$$

so the random variable $G(p_{i,n}X+1/n) - G(p_{i,n}\beta+1/n)$ is dominated by $MY_{i,n}$ and converges pointwisely to $Y_{i,n}$ as $n \to \infty$. Then, by Lebesgue's Dominated Convergence Theorem, we can generalize Lemma 1 to $G(x)$, and in turn generalize our bias bound.

Lemma 2. *Regarding the asymptotic behavior of* $\varphi_{i,n}$, *we have*

(a) $\lim\limits_{np_{i,n}\to 0} \dfrac{n^{2\lambda}\cdot\varphi_{i,n}}{np_{i,n}} = \beta^{1-2\lambda}$.

(b) $n^{-1+2\lambda}\ln n \cdot \varphi_{i,n} \le \beta^{1-2\lambda}/i$.

(c) *For any* $\delta > 0$, *there exists* M_δ *such that* $n^{-1+2\lambda}\ln n \sum_{i=1}^{\frac{n}{\delta\ln n}} \varphi_{i,n} \le M_\delta$ *for all* n.

(d) *For any* $\delta > 0$, *we have* $\lim\limits_{n\to\infty} n^{-1+2\lambda}\ln n \sum_{\frac{n}{\delta\ln n}\le i}^{n} \varphi_{i,n} = \infty$.

Lemma 1 is derived from Assumption (B) and the condition $\mathbb{E}\big[F(X)^2\big] < \infty$. Lemma 2 is derived from Assumption (A). The proofs are found in Appendix A.

2.4 Why is $\lambda < 0.5$ important?

As we note in Remark 4, the condition $\lambda < 0.5$ is necessary for the existence of $\mathbb{E}\big[F(X)^2\big]$. This existence is important because, briefly speaking, the Law of Large Numbers only holds when expected values exist. More precisely, we use the following lemma to prove convergence in probability in Theorem 1; if $\mathbb{E}\big[F(X)^2\big] = \infty$, the required uniform integrability is not satisfied, and the weighted averages of random variables we consider will not converge.

Lemma 3. *Assume the set of random variables $\big\{U_{i,n}/\varphi_{i,n}\big\}$ is uniformly integrable and $U_{i,n}$ $(1 \le i \le n)$ are independent for each fixed n. Assume $\lim_{np_{i,n}\to 0} \mathbb{E}[U_{i,n}]/\varphi_{i,n} = \ell$. Then,*

$$\lim_{n\to\infty} \frac{\sum_{i=1}^n U_{i,n}}{\sum_{i=1}^n \varphi_{i,n}} = \ell \quad \textit{in probability.}$$

Proof. This lemma is a combination of the Law of Large Numbers and the Stolz-Cesàro Theorem. We prove it in two steps.

First step, we prove

$$\lim_{n\to\infty} \frac{\sum_{i=1}^n U_{i,n} - \mathbb{E}[U_{i,n}]}{\sum_{i=1}^n \varphi_{i,n}} = 0 \quad \text{in probability.}$$

This is a generalized version of the Law of Large Numbers, saying that the weighted average of $U_{i,n}$ converges in probability to the weighted average of $\mathbb{E}[U_{i,n}]$. Since $\big\{U_{i,n}/\varphi_{i,n}\big\}$ is uniformly integrable, for any $\varepsilon > 0$ there exists N such that $\mathbb{E}\big[|U_{i,n}|I_{|U_{i,n}|>N\varphi_{i,n}}\big] < \varepsilon^2\varphi_{i,n}$ for all i, n. Our strategy is to divide the following into two parts,

$$\frac{\sum_{i=1}^n U_{i,n}}{\sum_{i=1}^n \varphi_{i,n}} = \frac{\sum_{i=1}^n U_{i,n}I_{|U_{i,n}|\le N\varphi_{i,n}}}{\sum_{i=1}^n \varphi_{i,n}} + \frac{\sum_{i=1}^n U_{i,n}I_{|U_{i,n}|> N\varphi_{i,n}}}{\sum_{i=1}^n \varphi_{i,n}},$$

and show that each part is close to its expectation. For the $|U_{i,n}| > N\varphi_{i,n}$ part, we have

$$\left|\mathbb{E}\left[\frac{\sum_{i=1}^n U_{i,n}I_{|U_{i,n}|>N\varphi_{i,n}}}{\sum_{i=1}^n \varphi_{i,n}}\right]\right| < \varepsilon^2 \quad \text{for all } n$$

by definition, so it has negligible expectation and can be bounded by Markov's Inequality:

$$\mathbb{P}\left(\left|\frac{\sum_{i=1}^n U_{i,n}I_{|U_{i,n}|>N\varphi_{i,n}}}{\sum_{i=1}^n \varphi_{i,n}} - \mathbb{E}\left[\frac{\sum_{i=1}^n U_{i,n}I_{|U_{i,n}|>N\varphi_{i,n}}}{\sum_{i=1}^n \varphi_{i,n}}\right]\right| > \varepsilon\right) < 2\varepsilon \quad \text{for all } n.$$

On the other hand, for the $|U_{i,n}| \le N\varphi_{i,n}$ part we have $\mathrm{Var}\big[U_{i,n}I_{|U_{i,n}|\le N\varphi_{i,n}}\big] \le N\varphi_{i,n}\mathbb{E}\big[|U_{i,n}|\big]$; and $\big\{U_{i,n}/\varphi_{i,n}\big\}$ being uniformly integrable implies $\mathbb{E}\big[|U_{i,n}|\big] \le M\varphi_{i,n}$ for some M, so

$$\mathrm{Var}\left[\frac{\sum_{i=1}^n U_{i,n}I_{|U_{i,n}|\le N\varphi_{i,n}}}{\sum_{i=1}^n \varphi_{i,n}}\right] \le \frac{\sum_{i=1}^n \varphi_{i,n}^2}{\big(\sum_{i=1}^n \varphi_{i,n}\big)^2} NM.$$

By Lemma 2(b), we have

$$\big(n^{-1+2\lambda}\ln n\big)^2 \sum_{i=1}^n \varphi_{i,n}^2 \le \beta^{2-4\lambda} \sum_{i=1}^\infty i^{-2} < \infty.$$

In contrast, by Lemma 2(d) we have

$$\lim_{n\to\infty}\left(n^{-1+2\lambda}\ln n\right)^2\Big(\sum_{i=1}^{n}\varphi_{i,n}\Big)^2=\infty.$$

Therefore,

$$\lim_{n\to\infty}\frac{\sum_{i=1}^{n}\varphi_{i,n}^2}{\left(\sum_{i=1}^{n}\varphi_{i,n}\right)^2}=0\quad\text{and}\quad\lim_{n\to\infty}\operatorname{Var}\left[\frac{\sum_{i=1}^{n}U_{i,n}I_{|U_{i,n}|\le N\varphi_{i,n}}}{\sum_{i=1}^{n}\varphi_{i,n}}\right]=0.$$

Thus, the $|U_{i,n}|\le N\varphi_{i,n}$ part concentrates to its expectation by Chebyshev's Inequality. The first step is completed.

Second step, we prove

$$\lim_{n\to\infty}\frac{\sum_{i=1}^{n}\mathbb{E}[U_{i,n}]}{\sum_{i=1}^{n}\varphi_{i,n}}=\lim_{np_{i,n}\to 0}\frac{\mathbb{E}[U_{i,n}]}{\varphi_{i,n}}.$$

This is a generalized version of the Stolz-Cesàro Theorem, saying that the limit ratio of two series equals the limit ratio of corresponding terms. By definition and Equation (4), for any $\varepsilon>0$ there exists δ such that

$$\frac{n}{\delta\ln n}\le i\le n\Rightarrow\left|\frac{\mathbb{E}[U_{i,n}]}{\varphi_{i,n}}-\ell\right|<\varepsilon\quad\text{for all } n.$$

In addition, we can bound

$$\left|\frac{\mathbb{E}[U_{i,n}]}{\varphi_{i,n}}-\ell\right|\le M\quad\text{for } 1\le i\le\frac{n}{\delta\ln n}$$

because $\{U_{i,n}/\varphi_{i,n}\}$ is uniformly integrable. The ratio $\big(\sum_{i=1}^{n}\mathbb{E}[U_{i,n}]\big)/\big(\sum_{i=1}^{n}\varphi_{i,n}\big)$ can be viewed as a weighted average of two parts, one from indices $\frac{n}{\delta\ln n}\le i\le n$ and the other from $1\le i<\frac{n}{\delta\ln n}$. By Lemma 2(d), the weight for the first part

$$\lim_{n\to\infty}n^{-1+2\lambda}\ln n\sum_{\frac{n}{\delta\ln n}\le i}^{n}\varphi_{i,n}=\infty;$$

whereas by Lemma 2(c), the weight for the second part

$$n^{-1+2\lambda}\ln n\sum_{i=1}^{\frac{n}{\delta\ln n}}\varphi_{i,n}<\infty.$$

Therefore, the first part dominates, so $\lim_{n\to\infty}\left|\big(\sum_{i=1}^{n}\mathbb{E}[U_{i,n}]\big)/\big(\sum_{i=1}^{n}\varphi_{i,n}\big)-\ell\right|<\varepsilon$. □

Combining Lemma 1(a)(c)(d) and Lemma 3, we immediately obtain the following.

Corollary 4. *For $\Upsilon:=\{ST\}$, $S/T\backslash S$ or $T/S\backslash T$, we have*

$$\lim_{n\to\infty}\frac{\sum_{i=1}^{n}\Big(F(p_{i,n}^{\Upsilon}+1/n)-\mathbb{E}\big[F(p_{i,n}^{\Upsilon}+1/n)\big]\Big)^2}{\eta\sum_{i=1}^{n}\varphi_{i,n}}=1\quad\textit{in probability.}$$

Now, we can asymptotically derive the normalization of a_n^{Υ}, $b_{i,n}$ and c_n, as defined in Section 2.2. A by-product is that the norms of natural phrase vectors converge to 1.

Theorem 2. *If we put* $a_n^{\{st\}} := 0$ *for all* $\{st\}$, *and set* $b_{i,n} := \mathbb{E}\big[F(p_{i,n}^{\{ST\}}+1/n)\big]$, $c_n := \big(\eta\sum_{i=1}^n \varphi_{i,n}\big)^{-1/2}$, *then*

$$\lim_{n\to\infty}\frac{1}{|\Lambda_n|}\sum_{\{st\}\in\Lambda_n}\mathbf{w}_n^{\{st\}}=\mathbf{0},\quad \lim_{n\to\infty}\frac{1}{n}\sum_{i=1}^n w_{i,n}^{\{ST\}}=0\quad \textit{and}\quad \lim_{n\to\infty}\|\mathbf{w}_n^{\{ST\}}\|=1$$

in probability.

Proof. By definition,

$$w_{i,n}^{\{st\}} := c_n\big(F(p_{i,n}^{\{st\}}+1/n)-a_n^{\{st\}}-b_{i,n}\big)=\frac{F(p_{i,n}^{\{st\}}+1/n)-\mathbb{E}\big[F(p_{i,n}^{\{ST\}}+1/n)\big]}{\sqrt{\eta\sum_{i=1}^n\varphi_{i,n}}}.$$

So Corollary 4 implies $\lim_{n\to\infty}\|\mathbf{w}_n^{\{ST\}}\|=1$.

Next, Lemma 1(c) implies that there exists M such that

$$\mathrm{Var}\big[F(p_{i,n}^{\{ST\}}+1/n)\big]\le M\varphi_{i,n}\quad\text{for all } i,n,$$

hence

$$\mathbb{E}\left[\left(\frac{1}{n}\sum_{i=1}^n w_{i,n}^{\{ST\}}\right)^2\right]=\frac{1}{n^2}\frac{\sum_{i=1}^n\mathrm{Var}\big[F(p_{i,n}^{\{ST\}}+1/n)\big]}{\eta\sum_{i=1}^n\varphi_{i,n}}\le\frac{1}{n^2}\frac{M}{\eta}\to 0\quad(\text{when } n\to\infty).$$

Therefore, by Chebyshev's Inequality we have $\lim_{n\to\infty}\frac{1}{n}\sum_{i=1}^n w_{i,n}^{\{ST\}}=0$ in probability.

Finally, the ordinary version of Law of Large Numbers implies that

$$\lim_{n\to\infty}\frac{1}{|\Lambda_n|}\sum_{\{st\}\in\Lambda_n}\frac{F(p_{i,n}^{\{st\}}+1/n)-\mathbb{E}\big[F(p_{i,n}^{\{ST\}}+1/n)\big]}{\sqrt{\mathrm{Var}\big[F(p_{i,n}^{\{ST\}}+1/n)\big]}}=0\quad\text{in probability;}$$

and since

$$\sqrt{\mathrm{Var}\big[F(p_{i,n}^{\{ST\}}+1/n)\big]}\le\sqrt{M\varphi_{i,n}}\le\sqrt{M\textstyle\sum_{i=1}^n\varphi_{i,n}},$$

we immediately have all entries of $\lim_{n\to\infty}\frac{1}{|\Lambda_n|}\sum_{\{st\}\in\Lambda_n}\mathbf{w}_n^{\{st\}}$ converge to 0. □

Therefore, if we set $a_n^{\{st\}}$, $b_{i,n}$ and c_n as in Theorem 2, all conditions in Definition 7, Definition 8 and Definition 9 are asymptotically satisfied. In addition, we have obtained the result stated in Claim 2.

In view of Corollary 4, if $\lambda < 0.5$ is not satisfied, the norms of natural phrase vectors will not converge. This prediction is experimentally verified in Section 5.2.

2.5 Proof of Theorem 1 and an Intuitive Explanation

Recall the equation (1) that $p_{i,n}^t$ is decomposed into a linear sum of $p_{i,n}^{s/t\backslash s}$ and $p_{i,n}^{\{st\}}$. The next lemma suggests that $F(p_{i,n}^t+1/n)$ can be treated like a linear sum of $F(p_{i,n}^{s/t\backslash s}+1/n)$ and $F(p_{i,n}^{\{st\}}+1/n)$ similarly, as if F has some linearity.

Lemma 5. *The set of random variables*

$$\left\{\left(F(p_{i,n}^T+1/n)-\pi_{S/T\backslash S}F(p_{i,n}^{S/T\backslash S}+1/n)-(1-\pi_{S/T\backslash S})F(p_{i,n}^{\{ST\}}+1/n)\right)^2/\varphi_{i,n}\right\}$$

is uniformly integrable, and

$$\lim_{np_{i,n}\to 0}\frac{\mathbb{E}\left[\left(F(p_{i,n}^T+1/n)-\pi_{S/T\backslash S}F(p_{i,n}^{S/T\backslash S}+1/n)-(1-\pi_{S/T\backslash S})F(p_{i,n}^{\{ST\}}+1/n)\right)^2\right]}{\varphi_{i,n}}=0.$$

Proof. The intuition is that when $np_{i,n}\to 0$, the probability $p_{i,n}$ is a small value compared to $1/n$, and since $p_{i,n}^{\Upsilon}$ ($\Upsilon:=\{ST\}$ or $\Upsilon:=S/T\backslash S$) is in the same scale as $p_{i,n}$, we can use the linear approximation $F(p_{i,n}^{\Upsilon}+1/n)\approx F'(1/n)p_{i,n}^{\Upsilon}+F(1/n)$. Formal proof is as below.

For brevity, we set

$$P_1:=p_{i,n}^{S/T\backslash S},\quad P_2:=p_{i,n}^{\{ST\}},\quad \pi:=\pi_{S/T\backslash S},\quad \text{and}\quad FA(x):=F(x+1/n)-F(p_{i,n}\beta+1/n).$$

By Equation (1) we have $p_{i,n}^T=\pi P_1+(1-\pi)P_2$, so $FA(p_{i,n}^T)=FA(\pi P_1+(1-\pi)P_2)$ lies between $FA(P_1)$ and $FA(P_2)$. Therefore,

$$\begin{aligned}\left(FA(p_{i,n}^T)-\pi FA(P_1)-(1-\pi)FA(P_2)\right)^2&\\ \le\left(FA(P_1)-FA(P_2)\right)^2&\\ \le 4FA(P_1)^2I_{FA(P_1)^2\ge FA(P_2)^2}+4FA(P_2)^2I_{FA(P_2)^2\ge FA(P_1)^2}.&\end{aligned}$$

By Lemma 1(c), $\left\{FA(P_1)^2/\varphi_{i,n}\right\}$ and $\left\{FA(P_2)^2/\varphi_{i,n}\right\}$ are uniformly integrable. So for any $\varepsilon>0$, we have $\mathbb{E}\left[FA(P_1)^2I_{FA(P_1)^2>N\varphi_{i,n}}\right]<\varepsilon\varphi_{i,n}$ and $\mathbb{E}\left[FA(P_2)^2I_{FA(P_2)^2>N\varphi_{i,n}}\right]<\varepsilon\varphi_{i,n}$ for some N. Consider the condition

$$\mathcal{C}:=\text{“}\textit{Either } FA(P_1)^2>N\varphi_{i,n} \textit{ or } FA(P_2)^2>N\varphi_{i,n}\text{”},$$

which is weaker than “$\left(FA(p_{i,n}^T)-\pi FA(P_1)-(1-\pi)FA(P_2)\right)^2>4N\varphi_{i,n}$”, and we have

$$\begin{aligned}\mathbb{E}\left[\left(FA(p_{i,n}^T)-\pi FA(P_1)-(1-\pi)FA(P_2)\right)^2I_{\mathcal{C}}\right]&\\ \le\mathbb{E}\left[4FA(P_1)^2I_{FA(P_1)^2>N\varphi_{i,n}}+4FA(P_2)^2I_{FA(P_2)^2>N\varphi_{i,n}}\right]&\\ <8\varepsilon\varphi_{i,n}.&\end{aligned}\tag{5}$$

So $\left\{\left(FA(p_{i,n}^T)-\pi FA(P_1)-(1-\pi)FA(P_2)\right)^2/\varphi_{i,n}\right\}$ is uniformly integrable.

The previous arguments suggest that we can neglect the possibility of $\mathcal{C}$ being satisfied, because the result (5) is arbitrarily small. Next, we consider the complement of $\mathcal{C}$, namely

$$\neg\mathcal{C} := \text{“}\textit{Both } FA(P_1)^2 \le N\varphi_{i,n} \textit{ and } FA(P_2)^2 \le N\varphi_{i,n}\text{”}.$$

Under this condition, intuitively $FA(P_1)$ and $FA(P_2)$ are restricted to a small range such that a linear approximation of F becomes valid. More precisely, we show that

$$\lim_{np_{i,n}\to 0} \frac{\mathbb{E}\Big[\big(FA(p_{i,n}^T) - \pi FA(P_1) - (1-\pi)FA(P_2)\big)^2 I_{\neg\mathcal{C}}\Big]}{\varphi_{i,n}} = 0,$$

which will complete the proof. For brevity, we set

$$FB(x) := F(x+1) - F(1), \quad U_1 := FB(nP_1), \quad U_2 := FB(nP_2),$$

use H to denote the inverse function of FB:

$$H(FB(x)) = x,$$

and put

$$J(u_1, u_2; \pi) := FB(\pi H(u_1) + (1-\pi)H(u_2)) - \pi u_1 - (1-\pi)u_2.$$

Note that FB, H and J do not depend on n, i, S or T. By Lemma 2(a), we can replace $\varphi_{i,n}$ with $np_{i,n}\cdot n^{-2\lambda}$; and note that

$$n^\lambda FA(x) = F(nx+1) - F(np_{i,n}\beta+1) \to FB(nx) \quad (\text{when } np_{i,n}\to 0),$$

so we only need prove

$$\lim_{np_{i,n}\to 0} \frac{\mathbb{E}\Big[J(U_1, U_2; \pi)^2 I_{\mathcal{D}}\Big]}{np_{i,n}} = 0,$$

where $\mathcal{D}$ is the condition

$$\mathcal{D} := \text{“}\textit{Both } U_1^2 \le Nnp_{i,n} \textit{ and } U_2^2 \le Nnp_{i,n}\text{”}.$$

Now, since $\frac{\partial}{\partial u_1}J(0,0;\pi) = \frac{\partial}{\partial u_2}J(0,0;\pi) = 0$, we have

$$\lim_{u_1^2+u_2^2\to 0} \frac{J(u_1,u_2;\pi)^2}{u_1^2+u_2^2} = 0 \quad \text{uniformly on } 0\le\pi\le 1.$$

Thus, when $np_{i,n}\to 0$, we have

$$\mathbb{E}\left[\frac{J(U_1,U_2;\pi)^2 I_{\mathcal{D}}}{np_{i,n}}\right] = \mathbb{E}\left[\frac{J(U_1,U_2;\pi)^2 I_{\mathcal{D}}}{U_1^2+U_2^2}\cdot\frac{U_1^2+U_2^2}{np_{i,n}}\right] \le 2N\mathbb{E}\left[\frac{J(U_1,U_2;\pi)^2 I_{\mathcal{D}}}{U_1^2+U_2^2}\right] \to 0.$$

□

Now, we are ready to prove Theorem 1. An intuitive discussion is given after the proof.

Proof of Theorem 1. As in Theorem 2, we set $a_n^{\{st\}} := 0$ and $c_n := \big(\eta \sum_{i=1}^n \varphi_{i,n}\big)^{-1/2}$. If $a^t := 0$ for all t, one can calculate that $\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^n w_{i,n}^T = 0$ in probability, as in the proof of Theorem 2 and with the help of Lemma 5. Thus, we can set $a_n^t := 0$. Then, by definition

$$\big(\mathcal{B}_n^{\{ST\}}\big)^2 = \frac{\sum_{i=1}^n \Big(F(p_{i,n}^{\{ST\}}+1/n) - \frac{1}{2}\big(F(p_{i,n}^S+1/n) + F(p_{i,n}^T+1/n)\big)\Big)^2}{\eta \sum_{i=1}^n \varphi_{i,n}}.$$

Next, by Lemma 5, Lemma 3 and Triangle Inequality, we can replace $F(p_{i,n}^T+1/n)$ with

$$\pi_{S/T\backslash S} F(p_{i,n}^{S/T\backslash S}+1/n) + (1-\pi_{S/T\backslash S}) F(p_{i,n}^{\{ST\}}+1/n),$$

and replace $F(p_{i,n}^S+1/n)$ with

$$\pi_{T/S\backslash T} F(p_{i,n}^{T/S\backslash T}+1/n) + (1-\pi_{T/S\backslash T}) F(p_{i,n}^{\{ST\}}+1/n).$$

For brevity, we put $\pi_1 := \pi_{S/T\backslash S}$, $\pi_2 := \pi_{T/S\backslash T}$, $FA(x) := F(x+1/n) - F(p_{i,n}\beta+1/n)$ and

$$W_{i,n}^{\{ST\}} := FA(p_{i,n}^{\{ST\}}), \quad W_{i,n}^{S/T\backslash S} := FA(p_{i,n}^{S/T\backslash S}), \quad W_{i,n}^{T/S\backslash T} := FA(p_{i,n}^{T/S\backslash T}).$$

We use "$\doteq$" to denote asymptotic equality at the limit $n \to \infty$. So

$$\big(\mathcal{B}_n^{\{ST\}}\big)^2 \doteq \frac{\sum_{i=1}^n \big((\pi_1+\pi_2)W_{i,n}^{\{ST\}} - \pi_1 W_{i,n}^{S/T\backslash S} - \pi_2 W_{i,n}^{T/S\backslash T}\big)^2}{4\eta \sum_{i=1}^n \varphi_{i,n}}.$$

Again, by Lemma 1(a)(b), Lemma 3 and Triangle Inequality, we can replace $W_{i,n}^{\Upsilon}$ with $\widetilde{W}_{i,n}^{\Upsilon} := W_{i,n}^{\Upsilon} - \mathbb{E}[W_{i,n}^{\Upsilon}]$ (where $\Upsilon := \{ST\}$, $S/T\backslash S$ or $T/S\backslash T$) respectively. So

$$\begin{aligned}\big(\mathcal{B}_n^{\{ST\}}\big)^2 \doteq{}& (\pi_1+\pi_2)^2 \frac{\sum_{i=1}^n \big(\widetilde{W}_{i,n}^{\{ST\}}\big)^2}{4\eta \sum_{i=1}^n \varphi_{i,n}} + \pi_1^2 \frac{\sum_{i=1}^n \big(\widetilde{W}_{i,n}^{S/T\backslash S}\big)^2}{4\eta \sum_{i=1}^n \varphi_{i,n}} + \pi_2^2 \frac{\sum_{i=1}^n \big(\widetilde{W}_{i,n}^{T/S\backslash T}\big)^2}{4\eta \sum_{i=1}^n \varphi_{i,n}} \\ &- 2\pi_1(\pi_1+\pi_2) \frac{\sum_{i=1}^n \widetilde{W}_{i,n}^{S/T\backslash S} \widetilde{W}_{i,n}^{\{ST\}}}{4\eta \sum_{i=1}^n \varphi_{i,n}} - 2\pi_2(\pi_1+\pi_2) \frac{\sum_{i=1}^n \widetilde{W}_{i,n}^{T/S\backslash T} \widetilde{W}_{i,n}^{\{ST\}}}{4\eta \sum_{i=1}^n \varphi_{i,n}} \\ &+ 2\pi_1\pi_2 \frac{\sum_{i=1}^n \widetilde{W}_{i,n}^{S/T\backslash S} \widetilde{W}_{i,n}^{T/S\backslash T}}{4\eta \sum_{i=1}^n \varphi_{i,n}}.\end{aligned}$$

By Corollary 4, we have

$$\frac{\sum_{i=1}^n \big(\widetilde{W}_{i,n}^{\{ST\}}\big)^2}{4\eta \sum_{i=1}^n \varphi_{i,n}} \doteq \frac{1}{4}, \quad \frac{\sum_{i=1}^n \big(\widetilde{W}_{i,n}^{S/T\backslash S}\big)^2}{4\eta \sum_{i=1}^n \varphi_{i,n}} \doteq \frac{1}{4} \quad \text{and} \quad \frac{\sum_{i=1}^n \big(\widetilde{W}_{i,n}^{T/S\backslash T}\big)^2}{4\eta \sum_{i=1}^n \varphi_{i,n}} \doteq \frac{1}{4}. \tag{6}$$

By Assumption (C), we have $\mathbb{E}\big[\widetilde{W}_{i,n}^{S/T\backslash S} \widetilde{W}_{i,n}^{T/S\backslash T}\big] = 0$, so applying Lemma 3 we get

$$\frac{\sum_{i=1}^n \widetilde{W}_{i,n}^{S/T\backslash S} \widetilde{W}_{i,n}^{T/S\backslash T}}{4\eta \sum_{i=1}^n \varphi_{i,n}} \doteq \lim_{np_{i,n}\to 0} \frac{\mathbb{E}\big[\widetilde{W}_{i,n}^{S/T\backslash S} \widetilde{W}_{i,n}^{T/S\backslash T}\big]}{4\eta\varphi_{i,n}} = 0.$$

Also by Assumption (C), we have $\mathbb{E}\big[\widetilde{W}_{i,n}^{S/T\setminus S}\widetilde{W}_{i,n}^{\{ST\}}\big] \geq 0$ and $\mathbb{E}\big[\widetilde{W}_{i,n}^{T/S\setminus T}\widetilde{W}_{i,n}^{\{ST\}}\big] \geq 0$, so

$$\lim_{n\to\infty} \frac{\sum_{i=1}^n \widetilde{W}_{i,n}^{S/T\setminus S}\widetilde{W}_{i,n}^{\{ST\}}}{4\eta\sum_{i=1}^n \varphi_{i,n}} \geq 0 \quad \text{and} \quad \lim_{n\to\infty} \frac{\sum_{i=1}^n \widetilde{W}_{i,n}^{T/S\setminus T}\widetilde{W}_{i,n}^{\{ST\}}}{4\eta\sum_{i=1}^n \varphi_{i,n}} \geq 0.$$

Therefore, $\lim_{n\to\infty}\big(\mathcal{B}_n^{\{ST\}}\big)^2 \leq \frac{1}{4}\big((\pi_1+\pi_2)^2+\pi_1^2+\pi_2^2\big) = \frac{1}{2}(\pi_1^2+\pi_2^2+\pi_1\pi_2)$. □

Intuitively, it is as if the word vector $\mathbf{w}_n^T$ can be decomposed into two components:

$$\mathbf{w}_n^T \simeq \pi_1\big(\widetilde{W}_{i,n}^{S/T\setminus S}\big)_{1\leq i\leq n} + (1-\pi_1)\big(\widetilde{W}_{i,n}^{\{ST\}}\big)_{1\leq i\leq n},$$

where "$\simeq$" is a heuristic symbol and other notations are as in Proof of Theorem 1. This decomposition is in correspondence to Equation (1), coming from the fact that a word target T can be divided into two types, one neighboring word S and the other not. Similarly, it is as if $\mathbf{w}_n^S$ can be decomposed into

$$\mathbf{w}_n^S \simeq \pi_2\big(\widetilde{W}_{i,n}^{T/S\setminus T}\big)_{1\leq i\leq n} + (1-\pi_2)\big(\widetilde{W}_{i,n}^{\{ST\}}\big)_{1\leq i\leq n}.$$

Thus, by taking the average $\frac{1}{2}(\mathbf{w}_n^S+\mathbf{w}_n^T)$, one combines the $\big(\widetilde{W}_{i,n}^{\{ST\}}\big)_{1\leq i\leq n}$ terms into a $1-(\pi_1+\pi_2)/2$ multiple of $\mathbf{w}_n^{\{ST\}}$, which will cancel $\mathbf{w}_n^{\{ST\}}$ in the bias $\big(\mathcal{B}_n^{\{ST\}}\big)^2$ and result in a $\frac{1}{4}(\pi_1+\pi_2)^2$ term, as shown in Proof of Theorem 1; meanwhile, the $\big(\widetilde{W}_{i,n}^{S/T\setminus S}\big)_{1\leq i\leq n}$ and $\big(\widetilde{W}_{i,n}^{T/S\setminus T}\big)_{1\leq i\leq n}$ terms tend to cancel out each other because they are independent, which corresponds to the $\mathbb{E}\big[\widetilde{W}_{i,n}^{S/T\setminus S}\widetilde{W}_{i,n}^{T/S\setminus T}\big] = 0$ argument in the proof and results in a $\frac{1}{4}(\pi_1^2+\pi_2^2)$ term in the bias bound.

It is noteworthy that, using the same argument as in Proof of Theorem 1, one can derive an upper bound for the Euclidean distance $\|\mathbf{w}_n^{\{ST\}}-\mathbf{w}_n^T\|$ instead of $\mathcal{B}_n^{\{ST\}}$, as $\lim_{n\to\infty}\|\mathbf{w}_n^{\{ST\}}-\mathbf{w}_n^T\| \leq \sqrt{2}\pi_1$. This is because if π_1 is small, $\mathbf{w}_n^T$ itself is close to $\mathbf{w}_n^{\{ST\}}$. Then, by Triangle Inequality we have

$$\lim_{n\to\infty}\mathcal{B}_n^{\{ST\}} \leq \lim_{n\to\infty}\frac{1}{2}(\|\mathbf{w}_n^{\{ST\}}-\mathbf{w}_n^T\|+\|\mathbf{w}_n^{\{ST\}}-\mathbf{w}_n^S\|) \leq \frac{\sqrt{2}}{2}(\pi_1+\pi_2).$$

This upper bound reflects how close one might expect $\frac{1}{2}(\mathbf{w}_n^S+\mathbf{w}_n^T)$ and $\mathbf{w}_n^{\{ST\}}$ to be, judged from $\|\mathbf{w}_n^{\{ST\}}-\mathbf{w}_n^T\|$ and $\|\mathbf{w}_n^{\{ST\}}-\mathbf{w}_n^S\|$ individually. It is looser than our bias bound presented in Theorem 1, suggesting that additive composition can perform better than one might expect. The difference between $\frac{\sqrt{2}}{2}(\pi_1+\pi_2)$ and $\sqrt{\frac{1}{2}(\pi_1^2+\pi_2^2+\pi_1\pi_2)}$ is exactly because $\big(\widetilde{W}_{i,n}^{S/T\setminus S}\big)_{1\leq i\leq n}$ and $\big(\widetilde{W}_{i,n}^{T/S\setminus T}\big)_{1\leq i\leq n}$ are independent and cancel out each other.

In view of the above explanation for Proof of Theorem 1, the technical points are as follows. First, the heuristic "$\simeq$" is not strict; there is difference between $\widetilde{W}_{i,n}^{\Upsilon}$ and $W_{i,n}^{\Upsilon}$ due to the expected value, and there is difference between $F(p_{i,n}^T+1/n)$ and a linear combination of $F(p_{i,n}^{S/T\setminus S}+1/n)$ and $F(p_{i,n}^{\{ST\}}+1/n)$. However, by Lemma 1(b) the expected values converge to 0, and by Lemma 5 a linear approximation holds if one takes limit. So the first

issues are settled. Second, the most importantly, terms in Equation (6) have to converge to constants independent of S and T, otherwise they cannot be separated from π_1 and π_2 in the calculation of $\left(\mathcal{B}_n^{\{ST\}}\right)^2$. Equation (6) comes from Corollary 4, which comes from Lemma 3 and is deeply related to Assumptions (A)(B) and the condition $\lambda < 0.5$.

Insights brought by our theory lead to several applications. First, as we found that the power law tail of natural language data requires $\lambda < 0.5$ for constructing additive compositional vector representations, our theory provides important guidance for the empirical research of Distributional Semantics (Section 3.1). Second, as we found that the component shared by $\mathbf{w}_n^T$ and $\mathbf{w}_n^S$ survives the averaging operation (i.e. the $\left(\widetilde{W}_{i,n}^{\{ST\}}\right)_{1\leq i\leq n}$ term), whereas independent terms $\left(\widetilde{W}_{i,n}^{S/T\backslash S}\right)_{1\leq i\leq n}$ and $\left(\widetilde{W}_{i,n}^{T/S\backslash T}\right)_{1\leq i\leq n}$ cancel out each other, we come to the idea of harnessing additive composition by engineering what is shared by the summands. Then, for example, we can make additive composition aware of word order (Section 3.2). Third, as one can read from Lemma 2(c)(d) and the second step of Proof of Lemma 3, it is important to realize that the behavior of vector representations is dominated by entries at dimensions corresponding to low-frequency words, namely $w_{i,n}^{\Upsilon}$ for $\frac{n}{\delta \ln n} \leq i \leq n$. This understanding has impact on dimension reduction (Section 3.3).

2.6 Hierarchical Pitman-Yor Process

In Assumptions (A)(B) of Theorem 1 we have required several properties to be satisfied by the probability values $p_{i,n}$ and $p_{i,n}^{\Upsilon}$. Meanwhile, $p_{i,n}$ and $p_{i,n}^{\Upsilon}$ $(1 \leq i \leq n)$ define distributions which words can be generated from. This setting resembles a Bayesian model where priors of word distributions are specified.

Conversely, by the well-known de Finetti's Theorem, an exchangeable random sequence of words (i.e., given any sequence sample, all permutations of that sample occur with the same probability) can be seen as if the words are drawn i.i.d. from a conditioned word distribution, where the distribution itself is drawn from a prior. A widely studied example is the Pitman-Yor Process (Pitman and Yor, 1997; Pitman, 2006); in this section, we use the process to define a generative model, from which Assumptions (A)(B) can be derived.

Definition 11. A Pitman-Yor Process $PY(\alpha, \theta)$ $(0 < \alpha < 1, \theta > -\alpha)$ defines a prior for word distributions, which is the prior corresponding to the exchangeable random sequence generated by the following Chinese Restaurant Process:

1. First, generate a new word.
2. At each step, let $C(\varpi)$ be the count of word ϖ, and $C := \sum_{\varpi} C(\varpi)$ the total count; let N be the number of distinct words. Then:
   (2.1) Generate a new word with probability $\dfrac{\theta + \alpha N}{\theta + C}$.
   (2.2) Or, generate a new copy of an existing word ϖ, with probability $\dfrac{C(\varpi) - \alpha}{\theta + C}$.

Definition 12. In the above process $PY(\alpha, \theta)$, we define $p(\varpi) := \lim \dfrac{C(\varpi)}{C}$, where limit is taken at Step $\to \infty$. Fix a word index i such that $p(\varpi_i) \geq p(\varpi_{i+1})$. Put $p_i := p(\varpi_i)$.

Theorem 3. *For a sequence generated by $PY(\alpha, \theta)$, we have $\lim C/N^{1/\alpha} = Z$ for some Z.*

Proof. This is Theorem 3.8 in Pitman (2006). □

Theorem 4. *We have $\lim_{i\to\infty} p_i \cdot i^{1/\alpha}\Gamma(1-\alpha)^{1/\alpha} = Z$, where Z is the same as in Theorem 3.*

Proof. This is Lemma 3.11 in Pitman (2006). □

Theorem 4 shows that, if words are generated by a Pitman-Yor Process $PY(\alpha, \theta)$, then p_i has a power law tail of index α. It is in the same form as Assumption (A), and when $\alpha \approx 1$, it approximates the Zipf's Law.

For two sequences generated by $PY(\alpha, \theta)$, their corresponding Z as in Theorem 3 may differ, even though the two sequences are generated with the same hyper-parameters α and θ. Nevertheless, the limit always exists, and Z is known to follow a statistical distribution that can be viewed as characterizing the prior of the word distribution $\{p_i\}$. The distribution of $Z^{-\alpha}$ is explicitly given in Pitman (2006), Theorem 3.8:

$$-\,\mathrm{d}\mathbb{P}(x \le Z^{-\alpha}) = \frac{\Gamma(\theta+1)}{\Gamma(\theta/\alpha+1)} x^{\theta/\alpha} g_\alpha(x)\,\mathrm{d}x \quad (x>0), \tag{7}$$

where $g_\alpha(x)$ is the Mittag-Leffler density function:

$$g_\alpha(x) := \frac{1}{\pi\alpha}\sum_{k=0}^{\infty} \frac{(-1)^{k+1}}{k!}\Gamma(\alpha k+1)\sin(\pi\alpha k) x^{k-1}.$$

In this article, we only need the fact that $\lim_{x\to 0} x g_\alpha(x)$ is a nonzero constant.

Next, we consider the co-occurrence probability $p^\Upsilon(\varpi)$, conditioned on ϖ being in the context of a target Υ. One first notes that $p^\Upsilon(\varpi)$ is likely to be related to $p(\varpi)$; i.e., frequent words are likely to occur in every context, regardless of target. To model this intuition, the idea of Hierarchical Pitman-Yor Process (Teh, 2006) is to adapt $PY(\alpha, \theta)$ such that in each step, if a new word is to be generated, it is no longer generated brand new, but drawn from another Pitman-Yor Process instead. This second Pitman-Yor Process serves as a "reference" which controls how frequently a word is likely to occur. More precisely, a Hierarchical Pitman-Yor Process $HPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ generates sequences as follows.

Definition 13. In $HPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$, instead of generating words directly, one generates a "reference" at each step, where the reference can refer to new words or existing words. We use ϱ to denote a reference and ϖ^ϱ the word referred to by the reference.

1. First step, generate a new reference which refers to a new word.
2. At each step, let $C(\varrho)$ be the count of reference ϱ, and $C(\varpi) := \sum_{\varpi^\varrho=\varpi} C(\varrho)$ the count of all references referring to word ϖ; let $C := \sum_\varpi C(\varpi)$ be the total count, $N_r(\varpi)$ the number of distinct references referring to ϖ, and $N_r := \sum_\varpi N_r(\varpi)$ the total number of distinct references; finally, let N_w be the number of distinct words.

   (2.1) Generate a new reference referring to a new word, with probability

$$\frac{1}{\theta_1 + C} \cdot \frac{\theta_1 + \alpha_1 N_r}{\theta_2 + N_r} \cdot (\theta_2 + \alpha_2 N_w).$$

(2.2) Generate a new reference referring to an existing word ϖ, with probability

$$\frac{1}{\theta_1 + C} \cdot \frac{\theta_1 + \alpha_1 N_r}{\theta_2 + N_r} \cdot (N_r(\varpi) - \alpha_2).$$

(2.3) Or, generate a new copy of an existing reference ϱ, with probability

$$\frac{1}{\theta_1 + C} \cdot (C(\varrho) - \alpha_1).$$

It is easy to see from definition that $HPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ generates an exchangeable word sequence; and if we focus on distinct references (i.e., ignoring (2.3), consider $N_r(\varpi)$ as "the count of word ϖ" in the ordinary Pitman-Yor Process), then the process becomes $PY(\alpha_2, \theta_2)$. We assume this is the same process which defines word probability $p(\varpi)$, so

$$p(\varpi) = \lim \frac{N_r(\varpi)}{N_r};$$

and we define the conditional probability $p^{\Upsilon}(\varpi)$ as:

$$p^{\Upsilon}(\varpi) := \lim \frac{C(\varpi)}{C}.$$

Thus, $HPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ indeed connects $p^{\Upsilon}(\varpi)$ to $p(\varpi)$. This connection between word probability and conditioned word probability has been explored in Teh (2006); in which, it is used in an n-gram language model to connect the bigram probability $p(w|u)$ to unigram probability $p(w)$, for deriving a smoothing method.

Unfortunately, a precise analysis on the above $p^{\Upsilon}(\varpi)$ is beyond the reach of the authors; instead, we consider a slightly modified process which is much simpler for our purpose.

Definition 14. A *Modified Hierarchical Pitman-Yor Process* $MHPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ generates sequences as follows. Using the same notation as in Definition 13:

1. First step, generate a new reference which refers to a new word.
2. At each step:

(2.1) Generate a new reference referring to a new word, with probability

$$\frac{1}{D} \cdot (\theta_2 + \alpha_2 N_w).$$

(2.2) Generate a new reference referring to an existing word ϖ, with probability

$$\frac{1}{D} \cdot (N_r(\varpi) - \alpha_2).$$

(2.3) Or, generate a new copy of an existing reference ϱ, with probability

$$\frac{1}{D} \cdot \frac{N_r(\varpi^{\varrho}) - \alpha_2}{\theta_1 + \alpha_1 N_r(\varpi^{\varrho})} \cdot (C(\varrho) - \alpha_1).$$

In above, D is a normalization factor that makes the probability values sum to 1:

$$D := \theta_2 + \alpha_2 N_w + \sum_{\varpi} \frac{N_r(\varpi) - \alpha_2}{\theta_1 + \alpha_1 N_r(\varpi)} \cdot (C(\varpi) + \theta_1).$$

$MHPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ modifies $HPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ by canceling the $(\theta_1+\alpha_1 N_r)/(\theta_2+N_r)$ factor in (2.1) and (2.2), and scaling (2.3) by $(N_r(\varpi^\varrho) - \alpha_2)/(\theta_1 + \alpha_1 N_r(\varpi^\varrho))$ instead. It is noteworthy that, since $\lim N_r = \infty$ and $\lim N_r(\varpi^\varrho) = \infty$, we have

$$\lim \frac{\theta_1 + \alpha_1 N_r}{\theta_2 + N_r} = \alpha_1 \quad \text{and} \quad \lim \frac{N_r(\varpi^\varrho) - \alpha_2}{\theta_1 + \alpha_1 N_r(\varpi^\varrho)} = \frac{1}{\alpha_1}.$$

So the asymptotic behaviors of $MHPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ and $HPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ are similar.

A favorable property of $MHPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$ is that, like $HPY(\alpha_1, \theta_1; \alpha_2, \theta_2)$, it becomes $PY(\alpha_2, \theta_2)$ when one focuses on distinct references, so we have

$$p(\varpi) = \lim \frac{N_r(\varpi)}{N_r} \tag{8}$$

as before; besides, if restricted to a specific word ϖ (i.e., ignoring (2.1), only consider the references referring to ϖ, and regard references as "words", $C(\varrho)$ as "the count of word ϱ", and $C(\varpi) := \sum_{\varpi^\varrho = \varpi} C(\varrho)$ as "the total count", $N_r(\varpi)$ as "the number of distinct words" in the ordinary Pitman-Yor Process), then the process becomes $PY(\alpha_1, \theta_1)$. Thus, by Theorem 3 we have

$$\lim \frac{C(\varpi)}{N_r(\varpi)^{1/\alpha_1}} = Z_\varpi \quad \text{for some } Z_\varpi. \tag{9}$$

Therefore, combining (8) and (9) we have

$$p^\Upsilon(\varpi) := \lim \frac{C(\varpi)}{C} = p(\varpi)^{1/\alpha_1} Z_\varpi \lim \frac{N_r^{1/\alpha_1}}{C}.$$

So $p^\Upsilon(\varpi)/p(\varpi)^{1/\alpha_1}$ is a constant multiple of Z_ϖ, which follows a distribution specified in Equation (7); and it is easy to see that $\{Z_\varpi\}$ for different ϖ are mutually independent. Thus, we have obtained Assumption (B1).

As for Assumption (B2), we assume $\theta_1 = 1$ and derive the distribution of Z_ϖ from (7):

$$-\,\mathrm{d}\mathbb{P}(z \le Z_w) = \frac{\alpha_1}{\Gamma(1/\alpha_1 + 1)} \frac{z^{-\alpha_1} g_{\alpha_1}(z^{-\alpha_1})}{z^2}\,\mathrm{d}z \quad (z > 0).$$

Since $\lim_{x \to 0} x g_\alpha(x)$ is a nonzero constant, the above probability density function is of order $o(z^{-2})$ when $z \to \infty$, so the random variable Z_w has a power law tail of index 1. Thus, Assumption (B2) is approximately satisfied when $\alpha_1 \approx 1$ and $\theta_1 = 1$.

3. Applications

In this section, we demonstrate three applications of our theory.

3.1 The Choice of Function F

The condition $\lambda < 0.5$ specifies a nontrivial constraint on the function F. In Section 2.4 we have shown that this is a necessary condition for the norms of natural phrase vectors to converge. The convergence of norms is an outstanding property that might affect not only additive composition but also the composition ability of vector representations in general. Specifically, we note that $F(x) = \ln x$ when $\lambda = 0$, and $F(x) = \sqrt{x}$ when $\lambda = 0.5$. It is straightforward to predict that these functions might perform better in composition tasks than functions that have larger λ, such as $F(x) := x$ or $F(x) := x \ln x$. In Section 5.2, we show experiments that verify the necessity of $\lambda < 0.5$ for our bias bound to hold, and in Section 6 we show that F indeed drastically affects additive compositionality as judged by human annotators – while $F(x) := \ln x$ and $F(x) := \sqrt{x}$ perform similarly well, $F(x) := x$ and $F(x) := x \ln x$ are much worse.

Different settings of function F have been considered in previous research, and speculations have been made about the reason of semantic additivity of some of the vector representations. In Pennington et al. (2014), the authors noted that logarithm is a homomorphism from multiplication to addition, and used this property to justify $F(x) := \ln x$ for training semantically additive word vectors, based but on the unverified hypothesis that multiplications of co-occurrence probabilities have specialties in semantics. On the other hand, Lebret and Collobert (2014) proposed to use $F(x) := \sqrt{x}$, which is motivated by the Hellinger distance between two probability distributions, and reported its being better than $F(x) := x$. Stratos et al. (2015) proposed a similar but more general and better-motivated model, which attributed $F(x) := \sqrt{x}$ to an optimal choice that stabilizes the variances of Poisson random variables. Based on the assumption that co-occurrence counts are generated by a Poisson Process, the authors pointed out that $F(x) := \sqrt{x}$ may have the effect of stabilizing the *variance* in estimating word vectors. In contrast, our theory shows clearly that F affects the *bias* of additive composition, besides the variance. All in all, none of the previous research can explain why $F(x) := \ln x$ and $F(x) := \sqrt{x}$ are *both* good choices but $F(x) := x$ is not.

Intuitively, the condition $\lambda < 0.5$ requires that $F(x)$ decreases steeply as x tends to 0. The steep slope has the effect of "amplifying" the fluctuations of lower co-occurrence probabilities, and "suppressing" higher ones as a result. Formally, this can be read from Lemma 1, where we show that $\mathrm{Var}[F(p_{i,n}^{\Upsilon}+1/n)]$ scales with $\varphi_{i,n} = p_{i,n}\big(p_{i,n}+(\beta n)^{-1}\big)^{-1+2\lambda}$. When $\lambda < 0.5$, the factor $\big(p_{i,n}+(\beta n)^{-1}\big)^{-1+2\lambda}$ decreases as $p_{i,n}$ increases, and the decrease is faster when λ is smaller. Thus, in the vector representations we consider, higher co-occurrence probabilities are "suppressed" more when λ is smaller. However, in practice higher co-occurrence probabilities can be more precisely estimated, and they often correspond to important words such as syntax markers. Therefore, though our bias bound only requires $\lambda < 0.5$, empirically one might obtain better semantic representations by adopting a λ that is not too small.

3.2 Handling Word Order in Additive Composition

By considering the vector representation $\mathbf{w}_n^{\{st\}}$ we have ignored word order and conflated the phrases "s t" and "t s". Though the meanings of the two might be related somehow,

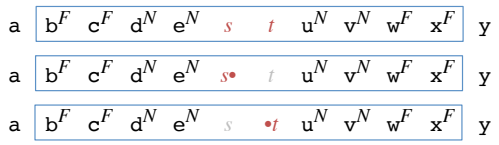

Figure 2: Surrounding the two-word phrase "s t", Near-far Contexts for $s\bullet$, $\bullet t$ and st are the same.

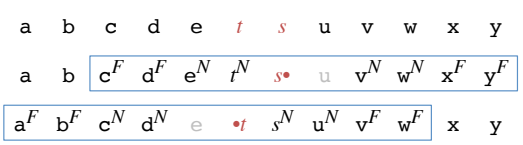

Figure 3: Surrounding the order-reversed phrase "t s", Near-far Contexts for $s\bullet$ and $\bullet t$ differ in their $\boldsymbol{N}$-$\boldsymbol{F}$ labels.

to treat a compositional framework as approximating $\mathbf{w}_n^{\{st\}}$ instead of $\mathbf{w}_n^{st}$ would certainly be troublesome, especially when one tries to extend our theory to longer phrases or even sentences. As the following famous example (Landauer et al., 1997) shows, meanings of sentences may differ greatly as word order changes.

a. *It was not the sales manager who hit the bottle that day, but the office worker with the serious drinking problem.*

b. *That day the office manager, who was drinking, hit the problem sales worker with a bottle, but it was not serious.*

Thus, it is necessary to handle the change of meanings brought by different word order. Traditionally, additive composition is considered unsuitable for this purpose, because one always has $\mathbf{w}_n^s + \mathbf{w}_n^t = \mathbf{w}_n^t + \mathbf{w}_n^s$. However, the commutativity can be broken by defining different contexts for "left-hand-side" words and "right-hand-side" words, denoted by $t\bullet$ and $\bullet t$, respectively. Then, the co-occurrence probabilities $p_{i,n}^{t\bullet}$ and $p_{i,n}^{\bullet t}$ will be different, so $\frac{1}{2}(\mathbf{w}_n^{s\bullet} + \mathbf{w}_n^{\bullet t})$ and $\frac{1}{2}(\mathbf{w}_n^{t\bullet} + \mathbf{w}_n^{\bullet s})$ are different vectors. In this section, we propose the *Near-far Context*, which specifies contexts for $s\bullet$ and $\bullet t$ such that the additive composition $\frac{1}{2}(\mathbf{w}_n^{s\bullet} + \mathbf{w}_n^{\bullet t})$ approximates the natural vector $\mathbf{w}_n^{st}$ of the *ordered* phrase "s t".

Definition 15. In *Near-far Context*, context words are assigned labels, either $\boldsymbol{N}$ or $\boldsymbol{F}$. For constructing vector representations, we use a lexicon of $\boldsymbol{N}$-$\boldsymbol{F}$ labeled words, and regard words with different labels as different entries in the lexicon. For any target, we label the nearer two words to each side by $\boldsymbol{N}$, and the farther two words to each side by $\boldsymbol{F}$. Except that, for the "left-hand-side" word $s\bullet$ we skip one word adjacent to the right; and similarly, for the "right-hand-side" word $\bullet t$ we skip one word adjacent to the left (Figure 2).

The idea behind Near-far Context is that, in the context of phrase "s t", each word is assigned an $\boldsymbol{N}$-$\boldsymbol{F}$ label the same as in the context of $s\bullet$ and $\bullet t$ (Figure 2). On the other hand, for targets s and t occurring in the order-reversed phrase "t s", context words are labeled differently for $s\bullet$ and $\bullet t$ (Figure 3). As we discussed in Section 2.2, the key fact about additive composition is that if a word token t comes from phrase "s t" or "t s", the context for this token of t is almost the same as the context of "s t" or "t s". By introducing different labels for context words of $t\bullet$ and $\bullet t$, we are able to distinguish "s t" from "t s". More precisely, as we discussed in Section 2.5, the component shared by $\mathbf{w}_n^{s\bullet}$ and $\mathbf{w}_n^{\bullet t}$ will survive in the average $\frac{1}{2}(\mathbf{w}_n^{s\bullet} + \mathbf{w}_n^{\bullet t})$, whereas non-shared components tend to cancel out each other. Thus, the additive composition $\frac{1}{2}(\mathbf{w}_n^{s\bullet} + \mathbf{w}_n^{\bullet t})$ will become closer to $\mathbf{w}_n^{st}$ rather than $\mathbf{w}_n^{ts}$, because $s\bullet$ and $\bullet t$ share context surrounding "s t" but not "t s".

Definition 16. Formally, as parallel to Definition 6, we define target $s\bullet\backslash t$ (resp. $s/\bullet t$) which counts every $s\bullet$ (resp. $\bullet t$) not at the left (resp. right) of word t (resp. s). And we use $\pi_{s\bullet\backslash t}$ (resp. $\pi_{s/\bullet t}$) to denote the probability of s (resp. t) not at the left (resp. right) of t (resp. s), conditioned on its occurrence. Practically, one can estimate $(1-\pi_{s\bullet\backslash t})$ (resp. $(1-\pi_{s/\bullet t})$) by the ratio of counts $C(st)/C(s)$ (resp. $C(st)/C(t)$). Then, we have equations

$$p^{s\bullet}_{i,n} = \pi_{s\bullet\backslash t} p^{s\bullet\backslash t}_{i,n} + (1-\pi_{s\bullet\backslash t}) p^{st}_{i,n} \quad \text{for all } i,n$$
$$p^{\bullet t}_{i,n} = \pi_{s/\bullet t} p^{s/\bullet t}_{i,n} + (1-\pi_{s/\bullet t}) p^{st}_{i,n} \quad \text{for all } i,n$$

as parallel to (1).

The following claim is parallel to Claim 1.

Claim 3. *Under conditions parallel to Claim 1, we have*

$$\lim_{n\to\infty} \mathcal{B}^{st}_n := \|\mathbf{w}^{st}_n - \frac{1}{2}(\mathbf{w}^{s\bullet}_n + \mathbf{w}^{\bullet t}_n)\| \le \sqrt{\frac{1}{2}(\pi^2_{s\bullet\backslash t} + \pi^2_{s/\bullet t} + \pi_{s\bullet\backslash t}\pi_{s/\bullet t})}.$$

In Section 5.3, we verify Claim 3 experimentally, and show that in contrast, the error $\|\mathbf{w}^{ts}_n - \frac{1}{2}(\mathbf{w}^{s\bullet}_n + \mathbf{w}^{\bullet t}_n)\|$ for approximating the order-reversed phrase "t s" can exceed this bias bound. Further, we demonstrate that by using the additive composition of Near-far Context vectors, one can indeed assess meaning similarities between ordered phrases.

3.3 Dimension Reduction

By far we have only discussed vector representations that have a high dimension equal to the lexicon size n. In practice, people mainly use low-dimensional "embeddings" of words to represent their meaning. Many of the embeddings, including SGNS and GloVe, can be formalized as linear dimension reduction, which is equivalent to the finding of a d-dimensional vector $\mathbf{v}^t$ (where $d \ll n$) for each target word t, and an (n,d)-matrix A such that $\sum_t L(A\mathbf{v}^t, \mathbf{w}^t_n)$ is minimized for some loss function $L(\cdot,\cdot)$. In other words, $A\mathbf{v}^t$ is trained as a good approximation for $\mathbf{w}^t_n$.

Naturally, we expect the loss function L to account for a crucial factor in word embeddings. Although there are empirical investigations on other detailed designs of embedding methods (e.g. how to count co-occurrences, see Levy et al. 2015), the loss functions have not been explicitly discussed previously. In this section, we discuss *how the loss functions would affect additive compositionality of word embeddings*, from a viewpoint of bounding the bias $\|\mathbf{v}^{\{st\}} - \frac{1}{2}(\mathbf{v}^s + \mathbf{v}^t)\|$.

SVD When L is the L^2-loss, its minimization has a closed-form solution given by the Singular Value Decomposition (SVD). More precisely, one considers a matrix whose j-th column is $\mathbf{w}^t_n$ where t is the j-th target word. Then, SVD factorizes the matrix into $U\Sigma V^\top$, where U, V are orthonormal and Σ is diagonal. Let Σ_d denote the truncated Σ to the top d singular values. Then, A is solved as $U\sqrt{\Sigma_d}$ and $\mathbf{v}^t$ the j-th column of $\sqrt{\Sigma_d}V^\top$. SVD has been used in Lebret and Collobert (2014), Stratos et al. (2015) and Levy et al. (2015).

For L^2-loss, we have

$$\|A\mathbf{v}^s - \mathbf{w}^s_n\| \le \varepsilon_1, \quad \|A\mathbf{v}^t - \mathbf{w}^t_n\| \le \varepsilon_2 \quad \text{and} \quad \|A\mathbf{v}^{\{st\}} - \mathbf{w}^{\{st\}}_n\| \le \varepsilon_3,$$

where ε_1, ε_2 and ε_3 are minimized. Thus, by Triangle Inequality we have

$$\|A \cdot \big(\mathbf{v}^{\{st\}} - \frac{1}{2}(\mathbf{v}^s + \mathbf{v}^t)\big)\| \leq \mathcal{B}_n^{\{st\}} + \frac{1}{2}(\varepsilon_1 + \varepsilon_2) + \varepsilon_3.$$

Further, by Claim 1 we can bound $\mathcal{B}_n^{\{st\}}$ for sufficiently large n, so $\|\mathbf{v}^{\{st\}} - \frac{1}{2}(\mathbf{v}^s + \mathbf{v}^t)\|$ is bounded in turn because A is a bounded operator. This bound suggests that word embeddings trained by SVD preserve additive compositionality.

However, the same argument does not directly apply to other loss functions because a general loss may not satisfy a triangle inequality, and a bound for Euclidean distance may not always transform to a bound for the loss, or vice versa. Specifically, we describe two widely used alternative embeddings in the following and discuss the effects of their loss.

GloVe The GloVe model (Pennington et al., 2014) trains a dimension reduction for vector representations with $F(x) := \ln x$. Let v_i^t be the i-th entry of $A\mathbf{v}^t$, and let C_i^t denote how many times target t co-occur with the i-th context word. Then, the loss function is

$$L(v_i^t, w_{i,n}^t) := f\big(C_i^t\big)(v_i^t - w_{i,n}^t)^2.$$

In words, GloVe uses a weighted L^2-loss and the weight is a function of the co-occurrence count C_i^t. The function f is set to constant when C_i^t is larger than a threshold, and decreases to 0 when $C_i^t \to 0$. To minimize the loss, GloVe uses stochastic gradient descent methods such as AdaGrad (Duchi et al., 2011).

SGNS The Skip-Gram with Negative Sampling (SGNS) model (Mikolov et al., 2013b) also trains a dimension reduction for vector representations with $F(x) := \ln x$. The training is based on the Noise Contrastive Estimation (NCE) (Gutmann and Hyvärinen, 2012), so the loss function inherits two parameters, namely the number k of noise samples per data point, and the noise distribution $p_{i,n}^{\text{noise}}$. The function is given below.

Claim 4. *Let v_i^t be the i-th entry of $A\mathbf{v}^t$. Then SGNS uses the loss function*

$$L(v_i^t, w_{i,n}^t) := C(t) D_{\phi_i}\big(v_i^t + \ln(kp_{i,n}^{\text{noise}}),\, w_{i,n}^t + \ln(kp_{i,n}^{\text{noise}})\big),$$

where $C(t)$ is the count of target t, and $D_\phi(\cdot,\cdot)$ is the Bregman divergence associated to the convex function

$$\phi(x) := \big(p_{i,n}^t + kp_{i,n}^{\text{noise}}\big) \ln\big(\exp(x) + kp_{i,n}^{\text{noise}}\big).$$

When $k \to +\infty$, D_ϕ converges to the Bregman divergence D_φ associated to $\varphi(x) := \exp(x)$.

Proof of Claim 4 is found in Appendix B. We draw a graph of the SGNS loss in Figure 4, where $D_\phi\big(v_i^t + \ln(kp_{i,n}^{\text{noise}}),\, w_{i,n}^t + \ln(kp_{i,n}^{\text{noise}})\big)$ is plotted on y-axis against $v_i^t - w_{i,n}^t$ on x-axis. Note that the graph grows faster at $x \to +\infty$ than $x \to -\infty$, suggesting that an overestimation of $w_{i,n}^t$ will be punished more than an underestimation. In addition, the loss function weighs more on high co-occurrence probabilities, as indicated by the $p_{i,n}^t$ coefficient in the equation of the limit curve (Figure 4). Thus, SGNS loss tends to enforce underestimation of $w_{i,n}^t$ for frequent context words (as overestimation is costly), and compensate $w_{i,n}^t$ for rare ones (i.e., overestimation on rare context words is affordable and will be done if necessary). This is a special property of SGNS which might have some smoothing effect.

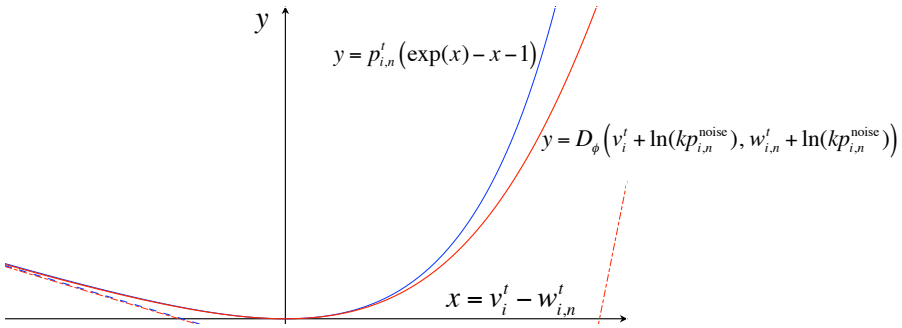

Figure 4: A graph of the SGNS loss function with two asymptotes (red), and its limit curve at $k \to +\infty$ with one asymptote (blue).

One common attribute shared by GloVe and SGNS but not SVD is that, the loss functions of GloVe and SGNS weigh less on rare context words. As a result, the trained $A\mathbf{v}^t$ may fail to precisely approximate the low co-occurrence part of $\mathbf{w}^t_n$. As we discussed in Section 2.5, entries corresponding to low-frequency words dominate the behavior of vector representations; thus, failing to precisely approximate this part might hinder the inheritance of additive compositionality from high-dimensional vector representations to low-dimensional embeddings. Therefore, we conjecture that word vectors trained by GloVe or SGNS might exhibit less additive compositionality compared to SVD, and the composition might be less respectful to our bias bound.

The above discussion is only exploratory and cannot fully comply with practice because, after $\mathbf{v}^t$ is trained by dimension reduction, people usually re-scale the norms of all $\mathbf{v}^t$ to 1, and then they use the normalized vectors in additive composition. It is not clear why this normalization step can usually result in better performance.

Nevertheless, in our experiments (Section 5.4), we find that word vectors trained by SVD preserve our bias bound well in additive composition, even after the normalization step is conducted. In contrast, vectors trained by GloVe or SGNS are less respectful to the bound. Further, in extrinsic evaluations (Section 6) we show that vectors trained by SVD can indeed be more additive compositional, as judged by human annotators.

4. Related Work

Additive composition is a classical approach to approximating meanings of phrases and/or sentences (Foltz et al., 1998; Landauer and Dutnais, 1997). Compared to other composition operations, vector addition/average has either served as a strong baseline (Mitchell and Lapata, 2008; Takase et al., 2016), or remained one of the most competitive methods until recently (Banea et al., 2014). Additive composition has also been successfully integrated into several NLP systems. For example, Tian et al. (2014) use vector additions for assessing semantic similarities between paraphrase candidates in a logic-based textual entailment recognition system (e.g. the similarity between "*blamed for death*" and "*cause loss of life*" is calculated by the cosine similarity between sums of word vectors $\mathbf{v}^{\text{blame}} + \mathbf{v}^{\text{death}}$ and $\mathbf{v}^{\text{cause}} + \mathbf{v}^{\text{loss}} + \mathbf{v}^{\text{life}}$); in Iyyer et al. (2015), average of vectors of words in a whole sentence/document is fed into a deep neural network for sentiment analysis and question answering, which

achieves near state-of-the-art performance with minimum training time. There are other semantic relations handled by vector additions as well, such as word analogy (e.g. the vector $\mathbf{v}^{\text{king}} - \mathbf{v}^{\text{man}} + \mathbf{v}^{\text{woman}}$ is close to $\mathbf{v}^{\text{queen}}$, suggesting "*man* is to *king* as *woman* is to *queen*", see Mikolov et al. 2013a), and synonymy (i.e. a set of synonyms can be represented by the sum of vectors of the words in the set, see Rothe and Schütze 2015). We expect all these utilities to be related to our theory of additive composition somehow, for example a link between additive composition and word analogy is hypothesized in Section 6.2. Ultimately, our theory would provide new insights into previous works, for instance, the insights about how to construct word vectors.

Lack of syntactic or word-order dependent effects on meaning is considered one of the most important issue of additive composition (Landauer, 2002). Driven by this point of view, a number of advanced compositional frameworks have been proposed to cope with word order and/or syntactic information (Mitchell and Lapata, 2008; Zanzotto et al., 2010; Baroni and Zamparelli, 2010; Coecke et al., 2010; Grefenstette and Sadrzadeh, 2011; Socher et al., 2012; Paperno et al., 2014; Hashimoto et al., 2014). The usual approach is to introduce new parameters that represent different word positions or syntactic roles. For example, given a two-word phrase, one can first transform the two word vectors by different matrices and then add the results, so the two matrices are parameters (Mitchell and Lapata, 2008); or, regarding different syntactic roles, one can assign matrices to adjectives and use them to modify vectors of nouns (Baroni and Zamparelli, 2010); further, one can insert neural network layers between parents and children in a syntactic tree (Socher et al., 2012). An empirical comparison of composition models can be found in Blacoe and Lapata (2012), with an accessible introduction to the literature. One theoretical issue of these methods, however, is the lack of learning guarantee. In contrast, our proposal of the Near-far Context demonstrates that word order can be handled within an additive compositional framework, being parameter-free and with a proven bias bound. Recently, Tian et al. (2016) further extended additive composition to realizing a formal semantics.

Error bounds in approximation schemes have been extensively studied in statistical learning theory (Vapnik, 1995; Gnecco and Sanguineti, 2008), and especially for neural networks (Niyogi and Girosi, 1999; Burger and Neubauer, 2001). Since we have formalized compositional frameworks as approximation schemes, there is a good chance to apply the theories of approximation error bounds to this problem, especially for advanced compositional frameworks that have many parameters. Though the theories are usually established on general settings, we see a great potential in using properties that are specific to natural language data, as we demonstrate in this work.

There have been consistent efforts toward understanding stochastic behaviors of natural language. Zipf's Law (Zipf, 1935) and its applications (Kobayashi, 2014), non-parametric Bayesian language models such as the Hierarchical Pitman-Yor Process (Teh, 2006), and the topic model (Blei, 2012) might further help refine our theory. For example, it can be fruitful to consider additive composition of topics.

5. Experimental Verification

In this section, we conduct experiments on the British National Corpus (BNC) (The BNC Consortium, 2007) to verify assumptions and predictions of our theory. The corpus contains

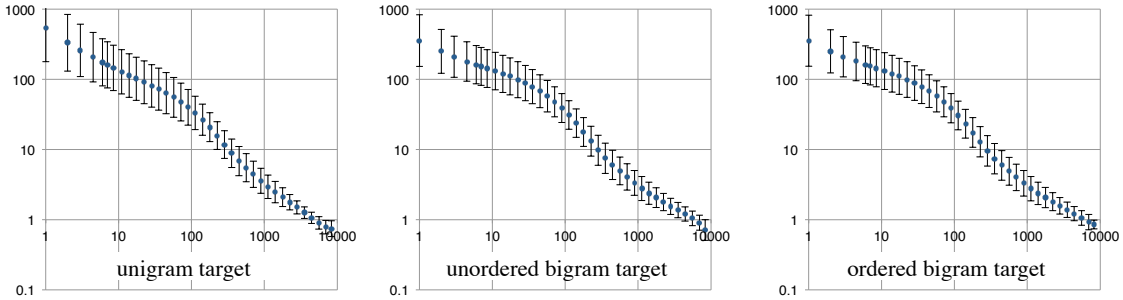

Figure 5: For each x coordinate, the log-log graphs show the average value of the x-th largest probability ratios $p^{\Upsilon}_{i,n}/p_{i,n}$ on y-axis. The ranking is taken among $1 \leq i \leq n$ with Υ fixed, and the average is taken across different Υ that are unigrams, unordered bigrams, or ordered bigrams respectively. Standard deviation is shown as error bar.

about 100M word tokens, including written texts and utterances in British English. For constructing vector representations we use lemmatized words annotated in the corpus, and for counting co-occurrences we use context windows that do not cross sentence boundaries. The size of the context windows is 5 to each side for a target word, and 4 for a target phrase. We extract all unigrams, ordered and unordered bigrams occurring more than 200 times as targets. This results in 16,210 unigrams, 45,793 ordered bigrams and 45,398 unordered bigrams. For the lexicon of context words we use the same set of unigrams.

5.1 Generalized Zipf's Law

Consider the probability ratio $p^{\Upsilon}_{i,n}/p_{i,n}$, where the target Υ can be a unigram, ordered bigram or unordered bigram. Assumption (B) of Theorem 1 states that $p^{\Upsilon}_{i,n}/p_{i,n}$ $(1 \leq i \leq n)$ can be viewed as independent sample points drawn from distributions that have a same power law tail of index 1. We verify this assumption in the following.

A power law distribution has two parameters, the index α and the lower bound m of the power law behavior. If a random variable X obeys a power law, the probability of $x \leq X$ conditioned on $m \leq X$ is given by

$$\mathbb{P}(x \leq X | m \leq X) = \frac{m^{\alpha}}{x^{\alpha}}. \tag{10}$$

For each fixed Υ, we estimate α and m from the sample $p^{\Upsilon}_{i,n}/p_{i,n}$ $(1 \leq i \leq n)$, using the method of Clauset et al. (2009). Namely, α is estimated by maximizing the likelihood of the sample, and m is sought to minimize the Kolmogorov-Smirnov statistic, which measures how well the theoretical distribution (10) fits the empirical distribution of the sample. After m is estimated, we plot all $p^{\Upsilon}_{i,n}/p_{i,n}$ greater than m in a log-log graph, against their ranking. If the sample points are drawn from a power law, the graph will be a straight line. Since Assumption (B) states that the power law tail is the same for all Υ and has index 1, we should obtain the same straight line for all Υ, and the slope of the line should be -1.

In Figure 5, we summarize the graphs described above for all Υ. More precisely, we plot the ranked probability ratios for each fixed Υ into the same log-log graph, and then show the average and standard deviation of the x-th largest probability ratios across different Υ.

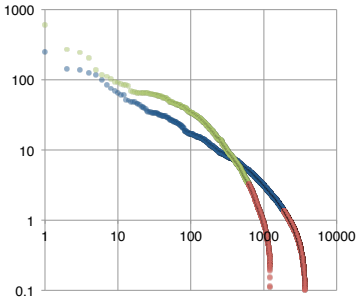

Figure 6: Log-log plot of probability ratios against ranks, for two fixed targets.

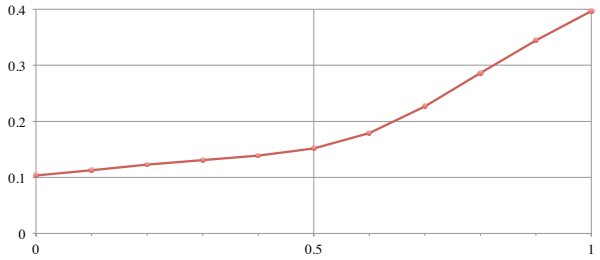

Figure 7: Plot of standard deviation of norms of natural phrase vectors against different λ.

The figure shows that for each target type, most data points lie within a narrow stripe of roughly the same shape, suggesting that the distribution of probability ratios for each fixed Υ is approximately the same. In addition, the shape can be roughly approximated by a straight line with slope -1, which suggests that the distribution is power law of index 1, verifying Assumption (B).

As a concrete example, in Figure 6 we show a log-log graph of the x-th largest probability ratios $p^s_{i,n}/p_{i,n}$ and $p^t_{i,n}/p_{i,n}$ $(1 \le i \le n)$, where s and t are two individual word targets. The red points are cut off because their y values are lower than the boundaries of power law behavior estimated from data. The blue and green points are the power law tails.

5.2 The Choice of Function F

In this section we experimentally verify the effects brought by different function F. Recall that F is parameterized by λ as defined in Definition 10. In Section 2.4, we have shown that $\mathbb{E}[F(X)^2] < \infty$ is a sufficient and necessary condition for the norms of natural phrase vectors to converge to 1. If X has a power law tail of index α, then the condition for $\mathbb{E}[F(X)^2] < \infty$ is $\lambda < \alpha/2$. So if we construct vector representations with different λ, only those vectors satisfying $\lambda < \alpha/2$ will have convergent norms. We verify this prediction first.

In Figure 7, we plot the standard deviation of the norms of natural phrase vectors on y-axis, against different λ values used for constructing the vectors. We tried $\lambda = 0, 0.1, \ldots, 1$. As the graph shows, as long as $\lambda < 0.5$, most of the norms lie within the range of 1 ± 0.1. In contrast, the observed standard deviation quickly explodes as λ gets larger. In addition, the transition point appears to be slightly larger than 0.5, which complies with the fact that the observed α is slightly larger than 1 (i.e., the slope $-1/\alpha$ of the power law tails in Figure 5 and Figure 6 appear to be slightly more gradual than -1).

Next, we investigate how F affects the Euclidean distance $\mathcal{B}^{\{st\}}_n$. In Figure 8, we plot

$$\mathcal{B}^{\{st\}}_n \text{ on } y\text{-axis}, \quad \text{against} \quad \sqrt{\frac{1}{2}(\pi^2_{s/t\backslash s} + \pi^2_{t/s\backslash t} + \pi_{s/t\backslash s}\pi_{t/s\backslash t})} \text{ on } x\text{-axis},$$

for every unordered bigram $\{st\}$. We tried four different choices of function F, as indicated above the graphs. For the choices (a) $F(p) := \ln p$ and (b) $F(p) := \sqrt{p}$, we verify the upper bound $y \le x$ as suggested by Claim 1. In contrast, the approximation errors seem no longer bounded when (c) $F(p) := p$ or (d) $F(p) := p \ln p$.

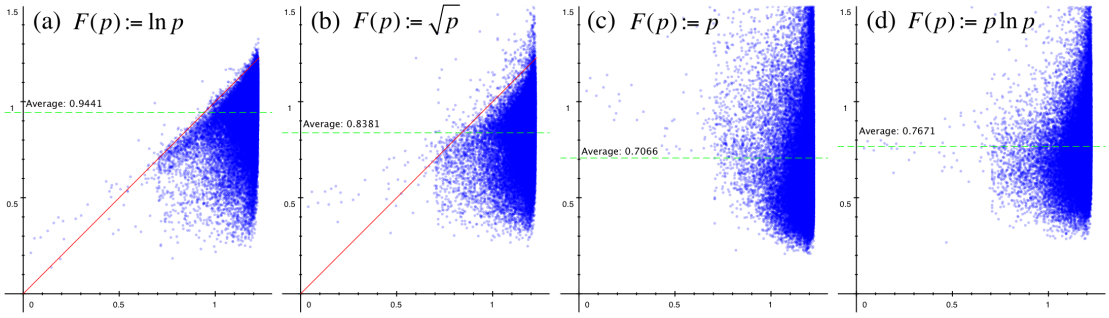

Figure 8: Approximation errors for unordered bigrams observed in BNC. The choice of F is shown above each graph. The theoretical upper bound $y \leq x$ is drawn as red solid lines in (a) and (b). The average error levels are drawn as green dashed lines.

In Section 6, we extrinsically evaluate the additive compositionality of vector representations, and find F a crucial factor there – while $F(p) := \ln p$ and $F(p) := \sqrt{p}$ evaluate similarly well, $F(p) := p$ and $F(p) := p \ln p$ do much worse. This suggests that our bias bound indeed has the power of predicting additive compositionality, demonstrating the usefulness of our theory. In contrast, it seems that the average level of approximation errors for observed bigrams (shown as green dashed lines in Figure 8) is less predictive, as the poor choices $F(p) := p$ and $F(p) := p \ln p$ actually have lower average error levels. This emphasizes a particular caveat that, choosing composition operations by minimizing the observed average error may not always be justifiable. Here if we consider the function F as a parameter in additive composition, and choose the one with the lowest average error observed, we will get the worst setting $F(p) := p$. Therefore, we see how important a learning theory for composition research is.

5.3 Handling Word Order in Additive Composition

For vector representations constructed from the Near-far Contexts (Section 3.2), we have a similar bias bound given by Claim 3. In this section, we experimentally verify the bound and qualitatively show that the additive composition of Near-far Context vectors can be used for assessing semantic similarities between *ordered* bigrams.

In Figure 9 and Figure 10, we plot

$$\text{(a) } \mathcal{B}_n^{st} \quad \text{and} \quad \text{(b) } \|\mathbf{w}_n^{ts} - \frac{1}{2}(\mathbf{w}_n^{s\bullet} + \mathbf{w}_n^{\bullet t})\| \quad \text{on } y\text{-axis,}$$

$$\text{against} \quad \sqrt{\frac{1}{2}(\pi_{s\bullet\backslash t}^2 + \pi_{s/\bullet t}^2 + \pi_{s\bullet\backslash t}\pi_{s/\bullet t})} \text{ on } x\text{-axis,}$$

for every ordered bigram st. We tried two settings of F, namely $F(p) := \ln p$ (Figure 9) and $F(p) := \sqrt{p}$ (Figure 10). In both cases, the approximation errors in (a) are bounded by $y \leq x$ (red solid lines) as suggested by Claim 3. In contrast, the approximation errors for order-reversed bigrams exceed this bound, showing that the additive composition of Near-far Context vectors actually recognizes word order.

In Table 4, we show the 8 nearest word pairs for each of 8 ordered bigrams, measured by cosine similarities between additive compositions of Near-far Context vectors. More

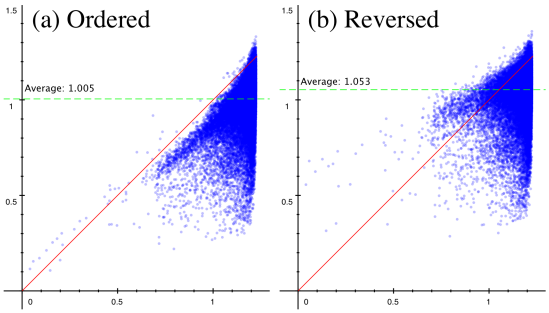

Figure 9: Near-far Context, $F(p) := \ln p$.

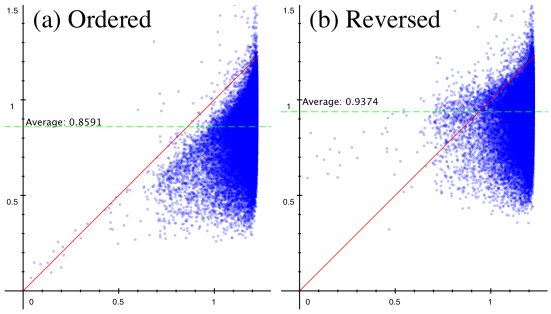

Figure 10: Near-far Context, $F(p) := \sqrt{p}$.

pose problem	***problem pose***	***tax rate***	***rate tax***
solve dilemma	*difficulty solve*	*income price*	income inflation
arise dilemma	*difficulty cause*	*income inflation*	*premium taxation*
solve difficulty	*difficulty tackle*	*taxation premium*	*premium inflation*
solve concern	*tendency solve*	*income premium*	*price income*
cause dilemma	*solution cause*	inflation income	taxation premium
tackle difficulty	*dilemma cause*	*taxation price*	*inflation income*
dilemma serious	*shortage solve*	premium taxation	*earnings taxation*
confront difficulty	*consequence solve*	*inflation premium*	*premium income*
high price	***price high***	***not enough***	***enough not***
low rate	*rate low*	*really sufficient*	*too never*
low premium	*level low*	*insufficient bother*	really never
low output	*value low*	*still bother*	*too really*
low value	*cost low*	*always want*	ought too
low cost	*premium low*	*always bother*	*too actually*
low wage	*output low*	*really prepared*	*too always*
low level	*inflation low*	*really unwilling*	*sufficient never*
low margin	*market low*	*really obliged*	*quite never*

Table 4: Top 8 similar word pairs, assessed by cosine similarities between additive compositions of Near-far Context vectors.

precisely, for word pairs "s_1 t_1" and "s_2 t_2", we calculate the cosine similarity between $\frac{1}{2}(\mathbf{v}^{s_1\bullet} + \mathbf{v}^{\bullet t_1})$ and $\frac{1}{2}(\mathbf{v}^{s_2\bullet} + \mathbf{v}^{\bullet t_2})$, where $\mathbf{v}^{s\bullet}$ and $\mathbf{v}^{\bullet t}$ are normalized 200-dimensional SVD reductions of $\mathbf{w}_n^{s\bullet}$ and $\mathbf{w}_n^{\bullet t}$, respectively, with $F(p) := \sqrt{p}$. The table shows that additive composition of Near-far Context vectors can indeed represent meanings of ordered bigrams, for example, "*pose problem*" is near to "*arise dilemma*" but not to "*dilemma arise*", and "*problem pose*" is near to "*difficulty cause*" but not to "*cause difficulty*". It is also noteworthy that "*not enough*" is similar to "*always want*", showing some degree of semantic compositionality beyond word level. We believe this ability of computing meanings of or-

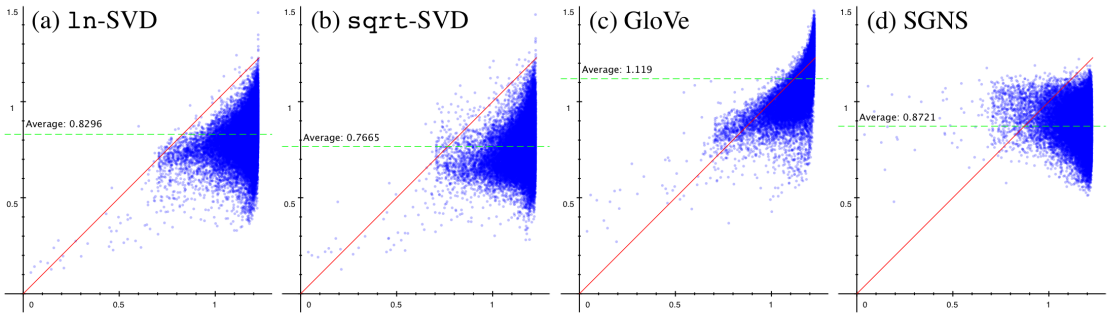

Figure 11: Approximation errors for different dimension reduction methods.

dered bigrams is already highly useful, because there are only a few bigrams whose meanings can be directly assessed from real corpora.

5.4 Dimension Reduction

In this section, we verify our prediction in Section 3.3 that vectors trained by SVD preserve our bias bound more faithfully than GloVe and SGNS. In Figure 11, we use *normalized* word vectors $\mathbf{v}^t$ that are constructed from the distributional vectors $\mathbf{w}_n^t$ by reducing to 200 dimensions using different reduction methods. We use SVD in (a) and (b), with $F(p) := \ln p$ in (a) and $F(p) := \sqrt{p}$ in (b). The GloVe model is shown in (c) and SGNS in (d), both of them using $F(p) := \ln p$. For each unordered bigram $\{st\}$ we plot

$$\|\mathbf{v}^{\{st\}} - \frac{1}{2}(\mathbf{v}^s + \mathbf{v}^t)\| \text{ on } y\text{-axis}, \quad \text{against} \quad \sqrt{\frac{1}{2}(\pi^2_{s/t\backslash s} + \pi^2_{t/s\backslash t} + \pi_{s/t\backslash s}\pi_{t/s\backslash t})} \text{ on } x\text{-axis}.$$

The graphs show that vectors trained by SVD still largely conform to our bias bound $y \leq x$ (red solid lines), but vectors trained by GloVe or SGNS no longer do. Our extrinsic evaluations in Section 6 also show that SVD might perform better than GloVe and SGNS.

6. Extrinsic Evaluation of Additive Compositionality

In this section, we test additive composition on human annotated data sets to see if our theoretical predictions correlate with human judgments. We conduct a phrase similarity task and a word analogy task.

6.1 Phrase Similarity

In a data set[3] created by Mitchell and Lapata (2010), phrase pairs are annotated with similarity scores. Each instance in the data is a (*phrase1*, *phrase2*, *similarity*) triplet, and each phrase consists of two words. The similarity score is annotated by humans, ranging from 1 to 7, indicating how similar the meanings of the two phrases are. For example, one annotator assessed the similarity between "*vast amount*" and "*large quantity*" as 7 (the highest), and the similarity between "*hear word*" and "*remember name*" as 1 (the lowest). Phrases are divided into three categories: Verb-Object, Compound Noun, and

3. `http://homepages.inf.ed.ac.uk/s0453356/`

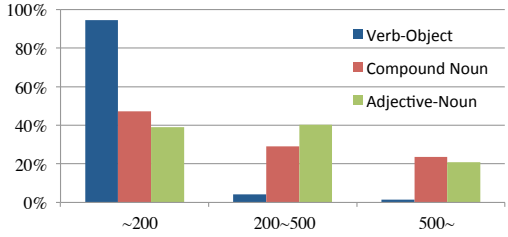

Figure 12: Distributions of how many times the phrases in the data occur as bigrams in BNC. The y-axis shows percentage and x-axis shows frequency range.

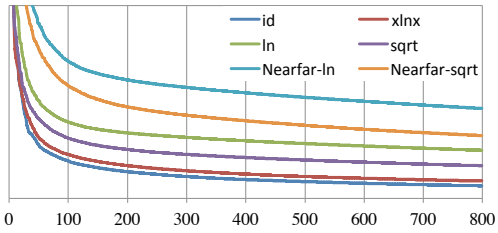

Figure 13: Top 800 singular values calculated by SVD. The y-axis shows singular value and x-axis shows rank. Different y-scales are used for different settings.

Adjective-Noun. Each category has 108 phrase pairs, and they are annotated by 18 human participants (i.e., 1,944 instances in each category). Using this data set, we can compare the human ranking of phrase similarity with the one calculated from cosine similarities between vector-based compositions. We use Spearman's ρ to measure how correlated the two rankings are.

Vector representations are constructed from BNC, with the same settings described in Section 5. We plot in Figure 12 the distributions of how many times the phrases in the data set occur as bigrams in BNC. The figure indicates that a large portion of the phrases are rare or unseen as bigrams, so their meanings cannot be directly assessed as natural vectors from the corpus. Therefore, the data is suitable for testing compositions of word vectors.

We reduce the high dimensional distributional word representations into 200-dimensional and normalize the vectors. The dimension 200 is selected by observing the top 800 singular values calculated by SVD. As illustrated in Figure 13, the decrease of singular values flattens to a constant rate at a rank of about 200. This suggests that the most characteristic features in the vector representations are projected into 200 dimensions. In our preliminary experiments, we have confirmed that a dimension of 200 performs better than 100, 500 or no dimension reduction.

For training word embeddings, we use the random projection algorithm (Halko et al., 2011) for SVD, and Stochastic Gradient Descent (SGD) (Bottou, 2012) for SGNS and GloVe. Since these are randomized algorithms, we run each test 20 times and report the mean performance with standard deviation. We tune SGD learning rates by checking convergence of the objectives, and get slightly better results than the default training parameters set in the software of SGNS[4] and GloVe[5].

As pointed out by Levy et al. (2015), there are other detailed settings that can vary in SGNS and GloVe. We make these settings close enough to be comparable but emphasize the differences of loss functions. More precisely, we use no subsampling and set the number of negative samples to 2 in SGNS, and use the default loss function in GloVe with the cutoff threshold set to 10. In addition, the default implementations of both SGNS and GloVe weigh context words by a function of distance to their targets, which we disable (i.e. equal

4. `https://code.google.com/p/word2vec/`
5. `http://nlp.stanford.edu/projects/glove/`

	Verb-Object	Compound Noun	Adjective-Noun
Ordinary-`id`-SVD	$.4029 \pm .0009$	$.4275 \pm .0009$	$.4160 \pm .0009$
Ordinary-`xlnx`-SVD	$.4204 \pm .0011$	$.4728 \pm .0013$	$.4511 \pm .0012$
Ordinary-`ln`-SVD	$\mathbf{.4369 \pm .0022}$	$\mathbf{.5187 \pm .0016}$	$.4604 \pm .0033$
Ordinary-`sqrt`-SVD	$.4318 \pm .0019$	$.5051 \pm .0020$	$.4790 \pm .0018$
Nearfar-`ln`-SVD	$.4204 \pm .0018$	$.5135 \pm .0020$	$.4491 \pm .0028$
Nearfar-`sqrt`-SVD	$\mathbf{.4359 \pm .0020}$	$\mathbf{.5193 \pm .0024}$	$.4873 \pm .0011$
SGNS	$.4273 \pm .0035$	$.4977 \pm .0025$	$\mathbf{.5125 \pm .0032}$
GloVe	$.4014 \pm .0046$	$.4986 \pm .0053$	$.4308 \pm .0062$
Tensor Product	$.4092 \pm .0033$	$.4801 \pm .0035$	$.4348 \pm .0048$
Upper Bound	.691	.693	.715
Muraoka et al.	.430	.481	.469
Deep Neural	.305	.385	.207

Table 5: Spearman's ρ in the phrase similarity task.

weights are used for all context words), so as to make it compatible with the contexts we consider in this article.

The test results are shown in Table 5. We compare different settings of function F, Ordinary and Near-far Contexts, and different dimension reductions. When using ordinary contexts and SVD reduction, we find that the functions `ln` ($F(p) := \ln p$) and `sqrt` ($F(p) := \sqrt{p}$) perform similarly well, whereas `id` ($F(p) := p$) and `xlnx` ($F(p) := p \ln p$) are much worse, confirming our predictions in Section 3.1. As for Near-far Context vectors (Section 3.2), we find that the Nearfar-`sqrt`-SVD setting has a high performance, demonstrating improvements brought by Near-far to additive composition. However, we note that Nearfar-`ln`-SVD is worse. One reason could be that the function `ln` emphasizes lower co-occurrence probabilities, which combined with Near-far labels could make the vectors more prone to data sparseness; or correlatively, some important syntactic markers might be obscured because they occur in high frequency. Finally, we note that SVD is consistently good and usually better than GloVe and SGNS, which supports our arguments in Section 3.3.

We report some additional test results for reference. In Table 5, the "Tensor Product" row shows the results of composing Ordinary-`ln`-SVD vectors by tensor product instead of average, which means that the similarity between two phrases "$s_1\ t_1$" and "$s_2\ t_2$" is assessed by taking product of the word cosine similarities $\cos(s_1, s_2) \cdot \cos(t_1, t_2)$. The numbers are worse than additive composition, suggesting that a similar phrase might be something more than a sequence of individually similar words. In the "Upper Bound" row, we show the best possible Spearman's ρ for this task, which are less than 1 because there are disagreements between human annotators. Compared to these numbers, we find that the performance of additive composition on compound nouns is remarkably high. Furthermore, in "Muraoka et al." we cite the best results reported in Muraoka et al. (2014), which has tested several compositional frameworks. In "Deep Neural", we also test additive composition of word vectors trained by deep neural networks (normalized 200-dimensional vectors trained by

Turian et al. 2010, using the model of Collobert et al. 2011). These results cannot be directly compared to each other because they construct vector representations from different corpora; but we can fairly say that additive composition is still a powerful method for assessing phrase similarity, and linear dimension reductions might be more suitable for training additively compositional word vectors than deep neural networks. Therefore, our theory on additive composition is about the state-of-the-art.

6.2 Word Analogy

Word analogy is the task of solving questions of the form "a is to b as c is to __?", and an elegant approach proposed by Mikolov et al. (2013a) is to find the word vector most similar to $\mathbf{v}^b - \mathbf{v}^a + \mathbf{v}^c$. For example, in order to answer the question "*man* is to *king* as *woman* is to __?", one needs to calculate $\mathbf{v}^{\text{king}} - \mathbf{v}^{\text{man}} + \mathbf{v}^{\text{woman}}$ and find out its most similar word vector, which will probably turn out to be $\mathbf{v}^{\text{queen}}$, indicating the correct answer *queen*.

As pointed out by Levy and Goldberg (2014a), the key to solving analogy questions is the ability to "add" (resp. "subtract") some aspects to (resp. from) a concept. For example, *king* is a concept of *human* that has the aspects of being *royal* and *male*. If we can "subtract" the aspect *male* from *king* and "add" the aspect *female* to it, then we will probably get the concept *queen*. Thus, the vector-based solution proposed by Mikolov et al. (2013a) is essentially assuming that "adding" and "subtracting" aspects can be realized by adding and subtracting word vectors. Why is this assumption admissible?

We believe this assumption is closely related to additive compositionality. Because, if an aspect is represented by an adjective (e.g. *male*) and a concept is represented by a noun (e.g. *human*), we can usually "add" the aspect to the concept by simply arranging the adjective and the noun to form a phrase (e.g. *male human*). Therefore, as the meaning of the phrase can be calculated by additive composition (e.g. $\mathbf{v}^{\text{male}} + \mathbf{v}^{\text{human}}$), we have indeed realized the "addition" of aspects by addition of word vectors. Specifically, since *man* $\approx$ *male human*, *king* $\approx$ *royal male human*, *woman* $\approx$ *female human* and *queen* $\approx$ *royal female human*, we expect the following by additive composition of phrases.

$$\begin{aligned}
\mathbf{v}^{\text{man}} &\approx \mathbf{v}^{\text{male}} + \mathbf{v}^{\text{human}} \\
\mathbf{v}^{\text{king}} &\approx \mathbf{v}^{\text{royal}} + \mathbf{v}^{\text{male}} + \mathbf{v}^{\text{human}} \\
\mathbf{v}^{\text{woman}} &\approx \mathbf{v}^{\text{female}} + \mathbf{v}^{\text{human}} \\
\mathbf{v}^{\text{queen}} &\approx \mathbf{v}^{\text{royal}} + \mathbf{v}^{\text{female}} + \mathbf{v}^{\text{human}}
\end{aligned}$$

Here, "$\approx$" denotes proximity between vectors in the sense of cosine similarity. From these approximate equations, we can imply that $\mathbf{v}^{\text{king}} - \mathbf{v}^{\text{man}} + \mathbf{v}^{\text{woman}} \approx \mathbf{v}^{\text{royal}} + \mathbf{v}^{\text{female}} + \mathbf{v}^{\text{human}} \approx \mathbf{v}^{\text{queen}}$, which solves the analogy question.

Therefore, we expect word analogy task to serve as an extrinsic evaluation for additive compositionality as well. For this reason, we conduct word analogy task on the standard Msr[6] (Mikolov et al., 2013a) and Google[7] (Mikolov et al., 2013b) data sets. Each instance in the data is a 4-tuple of words subject to "a is to b as c is to d", and the task is to find out d from a, b and c. We train word vectors with the same settings described in Section 5, but

6. http://research.microsoft.com/en-us/projects/rnn/
7. https://code.google.com/p/word2vec/

	id-SVD	xlnx-SVD	ln-SVD	sqrt-SVD	SGNS	GloVe
Google	$19.43 \pm .06$	$32.47 \pm .10$	$\mathbf{52.04 \pm .36}$	$51.28 \pm .25$	$45.16 \pm .44$	$48.39 \pm .44$
Msr	$17.36 \pm .06$	$33.85 \pm .12$	$\mathbf{66.67 \pm .26}$	$60.93 \pm .25$	$55.56 \pm .30$	$65.05 \pm .55$

Table 6: Accuracy (%) in the word analogy task.

using surface forms instead of lemmatized words in BNC. Tuples with out-of-vocabulary words are removed from data, which results in 4382 tuples in Msr and 8906 in Google[8].

The test results are shown in Table 6. Again, we find that ln and sqrt perform similarly well but id and xlnx are worse, confirming that the choice of function F can drastically affect performance on word analogy task as well, which we believe is related to additive compositionality. In addition, we confirm that SVD can perform better than SGNS and GloVe, which gives more support to our conjecture that vectors trained by SVD might be more compatible to additive composition.

7. Conclusion

In this article, we have developed a theory of additive composition regarding its bias. The theory has explained why and how additive composition works, making useful suggestions about improving additive compositionality, which include the choice of a transformation function, the awareness of word order, and the dimension reduction methods. Predictions made by our theory have been verified experimentally, and shown positive correlations with human judgments. In short, we have revealed the mechanism of additive composition.

However, we note that our theory is not "proof" of additive composition being a "good" compositional framework. As a generalization error bound usually is in machine learning theory, our bound for the bias does not show if additive composition is "good"; rather, it specifies some factors that can affect the errors. If we have generalization error bounds for other composition operations, a comparison between such bounds can bring useful insights into the choices of compositional frameworks in specific cases. We expect our bias bound to inspire more results in the research of semantic composition.

Moreover, we believe this line of theoretical research can be pursued further. In computational linguistics, the idea of treating semantics and semantic relations by algebraic operations on distributional context vectors is relatively new (Clarke, 2012). Therefore, the relation between linguistic theories and our approximation theory of semantic composition is left largely unexplored. For example, the intuitive distinction between compositional (e.g. *high price*) and non-compositional (e.g. *white lie*) phrases is currently ignored in our theory. Our bias bound treats both cases by a single collocation measure. Can one improve the bound by taking account of this distinction, and/or other kinds of linguistic knowledge? This is an intriguing question for future work.

8. These are about half the size of the original data sets.

Appendix A. Proof of Lemmas

In this appendix, we prove Lemma 1 and Lemma 2 in Section 2.3.

In order to prove Lemma 1, we first note the following equations:

$$\frac{F\big(x+(np_{i,n})^{-1}\big)-F\big(\beta+(np_{i,n})^{-1}\big)}{\big(1+(\beta np_{i,n})^{-1}\big)^{-1+\lambda}} \geq F(x)-F(\beta) \tag{11}$$

$$\frac{F\big(x+(np_{i,n})^{-1}\big)-F\big(\beta+(np_{i,n})^{-1}\big)}{\big(1+(\beta np_{i,n})^{-1}\big)^{\lambda}} = F\Big(\frac{x+(np_{i,n})^{-1}}{1+(\beta np_{i,n})^{-1}}\Big)-F(\beta) \tag{12}$$

Equation (11) can be obtained by analyzing the derivatives $F'(x)=x^{-1+\lambda}$, and Equation (12) immediately follows from the identity $z^{\lambda}(F(x)-F(y))=F(zx)-F(zy)$.

Proof of Lemma 1(a)(b). We calculate $e_{i,n}$ as follows. By definition,

$$e_{i,n}=\mathbb{E}\Big[\frac{Y_{i,n}}{\sqrt{\varphi_{i,n}}}I_{X\leq\beta}\Big]+\mathbb{E}\Big[\frac{Y_{i,n}}{\sqrt{\varphi_{i,n}}}I_{X\geq\beta}\Big].$$

Then, note that $Y_{i,n}\leq 0$ when $X\leq\beta$, and by (11) we have

$$0\geq\mathbb{E}\Big[\frac{Y_{i,n}}{\sqrt{\varphi_{i,n}}}I_{X\leq\beta}\Big]\geq\frac{\mathbb{E}\big[\big(F(X)-F(\beta)\big)I_{X\leq\beta}\big]}{\sqrt{1+(\beta np_{i,n})^{-1}}}. \tag{13}$$

From the condition $\mathbb{E}[F(X)^2]<\infty$ we have $\mathbb{E}[|F(X)|]<\infty$, so

$$(13)\geq-\mathbb{E}[|F(X)|]-|F(\beta)| \quad\text{and}\quad (13)\to 0 \text{ when } np_{i,n}\to 0. \tag{14}$$

Next, when $X\geq\beta$ we have $Y_{i,n}\geq 0$, and by (12) we get

$$0\leq\mathbb{E}\Big[\frac{Y_{i,n}}{\sqrt{\varphi_{i,n}}}I_{X\geq\beta}\Big]=\sqrt{1+(\beta np_{i,n})^{-1}}\cdot\mathbb{E}\left[\left(F\Big(\frac{X+(np_{i,n})^{-1}}{1+(\beta np_{i,n})^{-1}}\Big)-F(\beta)\right)I_{X\geq\beta}\right]. \tag{15}$$

By Assumption (B2), X obeys a power law at $X\geq\beta$, so if we put $Z:=X-\beta$, the probability density of Z is given by

$$-\,\mathrm{d}\mathbb{P}(z\leq Z)=\frac{\xi\,\mathrm{d}z}{(z+\beta)^2}\leq\frac{\xi\,\mathrm{d}z}{\sqrt{z}(z+\beta)^{3/2}}\quad(\text{where } z>0).$$

Thus,

$$(15)\leq\int_0^{\infty}\frac{F\Big(\dfrac{z}{1+(\beta np_{i,n})^{-1}}+\beta\Big)-F(\beta)}{\sqrt{\dfrac{z}{1+(\beta np_{i,n})^{-1}}}}\cdot\frac{\xi\,\mathrm{d}z}{(z+\beta)^{3/2}}. \tag{16}$$

The condition $\mathbb{E}[F(X)^2]<\infty$ implies $\lambda<0.5$, so $\dfrac{F(u+\beta)-F(\beta)}{\sqrt{u}}\to 0$ at $u\to\infty$. In addition the function is smooth on $[0,\infty)$, so it can be bounded by a constant M. Therefore,

$$(16)\leq\int_0^{\infty}M\cdot\frac{\xi\,\mathrm{d}z}{(z+\beta)^{3/2}}<\infty\quad\text{and}\quad(16)\to 0\text{ when }np_{i,n}\to 0. \tag{17}$$

The limit above is a consequence of Lebesgue's Dominated Convergence Theorem.

Combining (14) and (17), we have proven Lemma 1(a)(b). □

Proof of Lemma 1(c)(d). We calculate $\mathbb{E}[Y_{i,n}^2/\varphi_{i,n}]$ as follows. By definition,

$$\mathbb{E}[Y_{i,n}^2/\varphi_{i,n}] = \mathbb{E}\Big[\frac{Y_{i,n}^2}{\varphi_{i,n}} I_{X\le\beta}\Big] + \mathbb{E}\Big[\frac{Y_{i,n}^2}{\varphi_{i,n}} I_{X\ge\beta}\Big].$$

By (11) we have

$$\mathbb{E}\Big[\frac{Y_{i,n}^2}{\varphi_{i,n}} I_{X\le\beta}\Big] \le \frac{\mathbb{E}\big[\big(F(X)-F(\beta)\big)^2 I_{X\le\beta}\big]}{1+(\beta n p_{i,n})^{-1}}. \tag{18}$$

Then, from the condition $\mathbb{E}[F(X)^2]<\infty$, we have

$$(18) \text{ is uniformly integrable} \quad \text{and} \quad (18)\to 0 \text{ when } np_{i,n}\to 0. \tag{19}$$

Next, by (12) we get

$$\mathbb{E}\Big[\frac{Y_{i,n}^2}{\varphi_{i,n}} I_{X\ge\beta}\Big] = \big(1+(\beta n p_{i,n})^{-1}\big)\cdot \mathbb{E}\left[\left(F\Big(\frac{X+(np_{i,n})^{-1}}{1+(\beta np_{i,n})^{-1}}\Big)-F(\beta)\right)^2 I_{X\ge\beta}\right]. \tag{20}$$

By Assumption (B2), X obeys a power law at $X\ge\beta$, so if we put $Z := \dfrac{X-\beta}{1+(\beta np_{i,n})^{-1}}$, the probability density of Z is given by

$$\begin{aligned} -\,\mathrm{d}\mathbb{P}(z\le Z) &= \frac{1}{1+(\beta np_{i,n})^{-1}}\cdot\frac{\xi\,\mathrm{d}z}{\big(z+\dfrac{\beta}{1+(\beta np_{i,n})^{-1}}\big)^2} \quad (\text{where } z>0).\\ &\le \frac{1}{1+(\beta np_{i,n})^{-1}}\cdot\frac{\xi\,\mathrm{d}z}{z^2} \end{aligned}$$

Thus,

$$(20) \le \int_0^\infty \big(F(z+\beta)-F(\beta)\big)^2\cdot\frac{\xi\,\mathrm{d}z}{z^2}.$$

The condition $\mathbb{E}[F(X)^2]<\infty$ implies $\lambda<0.5$, so the above integral is finite. The integral is independent of i and n, so (20) is uniformly integrable; in addition, when $np_{i,n}\to 0$ we have $\big(z+\dfrac{\beta}{1+(\beta np_{i,n})^{-1}}\big)^2\to z^2$, so

$$\lim_{np_{i,n}\to 0}(20) = \int_0^\infty \big(F(z+\beta)-F(\beta)\big)^2\cdot\frac{\xi\,\mathrm{d}z}{z^2} \tag{21}$$

by Lebesgue's Dominated Convergence Theorem.

Combining (19) and (21), we have proven Lemma 1(c)(d). □

Proof of Lemma 2. Firstly, we note that

$$n^{-1+2\lambda}\cdot\varphi_{i,n} = \big(np_{i,n}+1/\beta\big)^{-1+2\lambda}\cdot p_{i,n}$$

by the definition of $\varphi_{i,n}$. So Lemma 2(a) immediately follows.

To prove Lemma 2(b), we simply apply Assumption (A) to the above.

To prove Lemma 2(c), we calculate

$$\begin{aligned}
n^{-1+2\lambda}\ln n\sum_{i=1}^{\frac{n}{\delta\ln n}}\varphi_{i,n} &= \sum_{i=1}^{\frac{n}{\delta\ln n}}\big(np_{i,n}+1/\beta\big)^{-1+2\lambda}\cdot p_{i,n}\ln n\\
&\le \sum_{i=1}^{\frac{n}{\delta\ln n}}(np_{i,n})^{-1+2\lambda}\cdot p_{i,n}\ln n\\
&\approx \big(\frac{n}{\ln n}\big)^{-1+2\lambda}\sum_{i=1}^{\frac{n}{\delta\ln n}} i^{-2\lambda} \qquad \text{(by Assumption (A))}\\
&\approx \big(\frac{n}{\ln n}\big)^{-1+2\lambda}\cdot\frac{1}{1-2\lambda}\big(\frac{n}{\delta\ln n}\big)^{1-2\lambda}\\
&= \frac{\delta^{-1+2\lambda}}{1-2\lambda},
\end{aligned}$$

so the required M_δ exists.

To prove Lemma 2(d), we note that by Equation (4),

$$\frac{n}{\delta\ln n}\le i\le n \Rightarrow np_{i,n}\le\delta.$$

Hence,

$$\begin{aligned}
n^{-1+2\lambda}\ln n\sum_{\frac{n}{\delta\ln n}\le i}^{n}\varphi_{i,n} &= \sum_{\frac{n}{\delta\ln n}\le i}^{n}\big(np_{i,n}+1/\beta\big)^{-1+2\lambda}\cdot p_{i,n}\ln n\\
&\ge \sum_{\frac{n}{\delta\ln n}\le i}^{n}(\delta+1/\beta)^{-1+2\lambda}\cdot p_{i,n}\ln n\\
&\approx (\delta+1/\beta)^{-1+2\lambda}\sum_{\frac{n}{\delta\ln n}\le i}^{n} i^{-1} \qquad \text{(by Assumption (A))}\\
&\approx (\delta+1/\beta)^{-1+2\lambda}\ln\big(\delta\ln n\big)\to\infty \quad \text{(when } n\to\infty\text{)}.
\end{aligned}$$

Therefore, Lemma 2(d) is proven. □

Appendix B. The Loss Function of SGNS

In this appendix, we discuss the loss function of SGNS. The model is originally proposed as an ad hoc objective function using the negative sampling technique (Mikolov et al., 2013b), without any explicit explanation on what is optimized and what is the loss. It is later shown that SGNS is a factorization of the shifted-PMI matrix (Levy and Goldberg, 2014b), but the loss function for this factorization remains unspecified. Here, we give a re-explanation of the SGNS model, with the loss function explicitly stated.

B.1 Noise Contrastive Estimation

The original objective function of SGNS is proposed as an adaptation of the Noise Contrastive Estimation (NCE) method, but in fact SGNS *is* using NCE without any adaptation. NCE (Gutmann and Hyvärinen, 2012) is a method for solving the classical problem that, given a sample $(x_i)_{i=1}^N$ (wherein $x_i \in \mathcal{X}$) drawn from an *unknown* probability distribution $\mathbb{P}_{\text{data}}$, and a function family $f(\cdot;\theta) : \mathcal{X} \to \mathbb{R}_{\geq 0}$ parameterized by θ, to find the optimal θ^* such that $f(x;\theta^*)$ approximates the distribution $\mathbb{P}_{\text{data}}(x)$ best. An alternative to NCE is the Maximum Likelihood Estimation (MLE), in which θ^* is chosen as to maximize the log-likelihood of the sample $(x_i)_{i=1}^N$, with respect to the constraint that $f(\cdot;\theta^*)$ should be a probability:

$$\theta^*_{\text{MLE}} = \arg\max_\theta \sum_{i=1}^N \ln f(x_i;\theta), \quad \text{where } \sum_{x\in\mathcal{X}} f(x;\theta) = 1.$$

For MLE, the constraint $\sum_{x\in\mathcal{X}} f(x;\theta) = 1$ is important, because $f(x;\theta)$ can tend to arbitrarily large if we maximize the log-likelihood without the constraint. NCE finds θ^* in a different way. It firstly mixes (x_i) with a noise sample drawn from a *known* distribution $\mathbb{P}_{\text{noise}}$, each data point x_i mixed with k noise points $y_{i,1}, \ldots, y_{i,k} \sim \mathbb{P}_{\text{noise}}$. Hence,

$$\mathbb{P}(x \text{ is data} \mid x) = \frac{\mathbb{P}_{\text{data}}(x)}{\mathbb{P}_{\text{data}}(x) + k\mathbb{P}_{\text{noise}}(x)}, \tag{22}$$

which gives the probability of a given point $x \in \mathcal{X}$ being a data point. $\mathbb{P}_{\text{data}}$ is unknown in (22), so we approximate $\mathbb{P}(x \text{ is data} \mid x)$ by $g(x;\theta)$ as below:

$$g(x;\theta) := \frac{f(x;\theta)}{f(x;\theta) + k\mathbb{P}_{\text{noise}}(x)}. \tag{23}$$

Then, NCE maximizes the log-likelihood of "x_i *being data and* $y_{i,1}, \ldots, y_{i,k}$ *being noise*":

$$\theta^*_{\text{NCE}} = \arg\max_\theta \sum_{i=1}^N \Big(\ln g(x_i;\theta) + \sum_{j=1}^k \ln(1 - g(y_{i,j};\theta)) \Big). \tag{24}$$

The most important point of NCE is that, $f(x;\theta)$ will *not* tend to infinity even we maximize (24) *without* the constraint $\sum_{x\in\mathcal{X}} f(x;\theta) = 1$. This is because making $f(x;\theta)$ large will accordingly make $1 - g(y_{i,j};\theta)$ small, which will *decrease* the likelihood of "$y_{i,1}, \ldots, y_{i,k}$ *being noise*". No longer necessary to repeatedly calculate $\sum_{x\in\mathcal{X}} f(x;\theta)$ during parameter update, NCE usually results in efficient training algorithms.

B.2 The Skip-Gram with Negative Sampling Model

Let p_i^t be the co-occurrence probability of the i-th word, conditioned on it being in the context of a target word t. SGNS approximates p_i^t by the function family

$$f(i,t;\mathbf{u},\mathbf{v}) := \exp(\mathbf{u}_i \cdot \mathbf{v}^t + \ln(kp_i^{\text{noise}})),$$

using NCE to optimize parameters. Here, $\mathbf{u}$ and $\mathbf{v}$ are parameters of the function family, whose columns are vectors $\mathbf{u}_i$ and $\mathbf{v}^t$, corresponding to the i-th context word and the word

target t, respectively. The training data $\mathcal{C}$ is a collection of co-occurring context-target word pairs. On the other hand, k and p_i^{noise} are constants in the definition of the function family, where k is the number of noise points drawn for each training instance, and p_i^{noise} is the probability value of the i-th context word being drawn from the noise distribution $\mathbb{P}_{\text{noise}}$.

Substituting the above $f(i, t; \mathbf{u}, \mathbf{v})$ into (23), we get

$$g(i, t; \mathbf{u}, \mathbf{v}) = \frac{\exp(\mathbf{u}_i \cdot \mathbf{v}^t + \ln(kp_i^{\text{noise}}))}{\exp(\mathbf{u}_i \cdot \mathbf{v}^t + \ln(kp_i^{\text{noise}})) + kp_i^{\text{noise}}} = \sigma(\mathbf{u}_i \cdot \mathbf{v}^t),$$

where $\sigma(x) = 1/\big(1 + \exp(-x)\big)$ is the sigmoid function.

Substituting the obtained $g(i, t; \mathbf{u}, \mathbf{v})$ into (24), we get

$$(\mathbf{u}, \mathbf{v})^*_{\text{NCE}} = \arg\max_{\mathbf{u},\mathbf{v}} \sum_{(i,t)\in\mathcal{C}} \Big(\ln \sigma(\mathbf{u}_i \cdot \mathbf{v}^t) + \sum_{y\sim\mathbb{P}_{\text{noise}}} \ln(1 - \sigma(\mathbf{u}_y \cdot \mathbf{v}^t))\Big)$$

which is exactly the objective function of SGNS proposed in Mikolov et al. (2013b).

Since SGNS is using $f(i, t; \mathbf{u}, \mathbf{v}) = \exp(\mathbf{u}_i \cdot \mathbf{v}^t + \ln(kp_i^{\text{noise}}))$ to approximate p_i^t, it is using $\mathbf{u}_i \cdot \mathbf{v}_t$ to approximate $w_i^t := \ln p_i^t - \ln(kp_i^{\text{noise}})$. This w_i^t is in the form of vector entries of distributional representations as given in Definition 4. Thus, SGNS can be viewed as a dimension reduction of the distributional representations we consider in this article.

Proof of Claim 4. To calculate its loss function, we consider the objective of SGNS:

$$O(\mathbf{u}, \mathbf{v}) := \sum_{(i,t)\in\mathcal{C}} \Big(\ln \sigma(\mathbf{u}_i \cdot \mathbf{v}^t) + \sum_{y\sim\mathbb{P}_{\text{noise}}} \ln(1 - \sigma(\mathbf{u}_y \cdot \mathbf{v}^t))\Big).$$

The sum is taken across all context-target pairs in $\mathcal{C}$. We regroup the summands by each distinct target, and note that conditioned on an occurrence of target t, the probability for one to encounter the i-th context word co-occurring in the training data is p_i^t, whereas the expected times for one to draw the context word from noise is given by kp_i^{noise}. So we have

$$O(\mathbf{u}, \mathbf{v}) = \sum_t C(t) \sum_i \Big(p_i^t \ln \sigma(\mathbf{u}_i \cdot \mathbf{v}^t) + kp_i^{\text{noise}} \ln(1 - \sigma(\mathbf{u}_n \cdot \mathbf{v}^t))\Big)$$

where $C(t)$ is the occurrence count of t. Now, we know that the optimal $O(\mathbf{u}, \mathbf{v})$ is achieved at $\mathbf{u}_i \cdot \mathbf{v}^t = w_i^t$, so we define

$$M := \sum_t C(t) \sum_i \Big(p_i^t \ln \sigma(w_i^t) + kp_i^{\text{noise}} \ln(1 - \sigma(w_i^t))\Big).$$

Then, to maximize $O(\mathbf{u}, \mathbf{v})$ is to minimize $M - O(\mathbf{u}, \mathbf{v})$, and by some calculation we obtain

$$M - O(\mathbf{u}, \mathbf{v}) = \sum_t C(t) \sum_i D_\phi\big(\mathbf{u}_i \cdot \mathbf{v}^t + \ln(kp_i^{\text{noise}}),\ w_i^t + \ln(kp_i^{\text{noise}})\big),$$

where $D_\phi(p,\, q) := \phi(p) - \phi(q) - \phi'(q)(p - q)$ is the Bregman divergence associated with the convex function

$$\phi(x) = (p_i^t + kp_i^{\text{noise}}) \ln(\exp(x) + kp_i^{\text{noise}}).$$

This is the loss function given in Claim 4. The limit of D_ϕ at $k \to +\infty$ is easily derived. □

References

Carmen Banea, Di Chen, Rada Mihalcea, Claire Cardie, and Janyce Wiebe. Simcompass: Using deep learning word embeddings to assess cross-level similarity. In *Proceedings of SemEval*, 2014.

Marco Baroni and Roberto Zamparelli. Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In *Proceedings of EMNLP*, 2010.

William Blacoe and Mirella Lapata. A comparison of vector-based representations for semantic composition. In *Proceedings of EMNLP*, 2012.

David M. Blei. Probabilistic topic models. *Commun. ACM*, 55(4), 2012.

Gemma Boleda, Marco Baroni, The Nghia Pham, and Louise McNally. Intensionality was only alleged: On adjective-noun composition in distributional semantics. In *Proceedings of IWCS*, 2013.

Léon Bottou. Stochastic gradient descent tricks. In Grégoire Montavon, Geneviève B. Orr, and Klaus-Robert Müller, editors, *Neural Networks: Tricks of the Trade*. Springer, Berlin, 2012.

Martin Burger and Andreas Neubauer. Error bounds for approximation with neural networks. *Journal of Approximation Theory*, 112(2), 2001.

Kenneth Ward Church and Patrick Hanks. Word association norms, mutual information, and lexicography. *Comput. Linguist.*, 16(1), 1990.

Daoud Clarke. A context-theoretic framework for compositionality in distributional semantics. *Comput. Linguist.*, 38(1), 2012.

Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Rev.*, 51(4), 2009.

Bob Coecke, Mehrnoosh Sadrzadeh, and Stephen Clark. Mathematical foundations for a compositional distributional model of meaning. *Linguistic Analysis*, 2010.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. Natural language processing (almost) from scratch. *J. Mach. Learn. Res.*, 12, 2011.

Á. Corral, G. Boleda, and R. Ferrer i Cancho. Zipf's law for word frequencies: Word forms versus lemmas in long texts. *PLoS ONE*, 10(7), 2015.

Ido Dagan, Fernando Pereira, and Lillian Lee. Similarity-based estimation of word cooccurrence probabilities. In *Proceedings of ACL*, 1994.

Georgiana Dinu, Nghia The Pham, and Marco Baroni. General estimation and evaluation of compositional distributional semantic models. In *Proceedings of the Workshop on Continuous Vector Space Models and their Compositionality*, 2013.

John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.*, 12, 2011.

Peter W. Foltz, Walter Kintsch, and Thomas K. Landauer. The measurement of textual coherence with latent semantic analysis. *Discourse Process*, 15, 1998.

Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural Comput.*, 4(1), 1992.

Giorgio Gnecco and Marcello Sanguineti. Approximation error bounds via rademachers complexity. *Applied Mathematical Sciences*, 2(4), 2008.

Edward Grefenstette and Mehrnoosh Sadrzadeh. Experimental support for a categorical compositional distributional model of meaning. In *Proceedings of EMNLP*, 2011.

Emiliano Guevara. A regression model of adjective-noun compositionality in distributional semantics. In *Proceedings of the Workshop on GEometrical Models of Natural Language Semantics*, 2010.

Michael U. Gutmann and Aapo Hyvärinen. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *J. Mach. Learn. Res.*, 13 (1), 2012.

Le Quan Ha, E. I. Sicilia-Garcia, Ji Ming, and F. J. Smith. Extension of zipf's law to words and phrases. In *Proceedings of Coling*, 2002.

N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53 (2), 2011.

Zellig S. Harris. Distributional structure. *Word*, 10, 1954.

Kazuma Hashimoto, Pontus Stenetorp, Makoto Miwa, and Yoshimasa Tsuruoka. Jointly learning word representations and composition functions using predicate-argument structures. In *Proceedings of EMNLP*, 2014.

Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of ACL*, 2015.

Hayato Kobayashi. Perplexity on reduced corpora. In *Proceedings of ACL*, 2014.

Thomas K. Landauer. On the computational basis of learning and cognition: Arguments from LSA. In N. Ross, editor, *The psychology of learning and motivation*, volume 41. Academic Press, 2002.

Thomas K Landauer and Susan T. Dutnais. A solution to platos problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review*, 104(2), 1997.

Thomas K. Landauer, Darrell Laham, Bob Rehder, and M. E. Schreiner. How well can passage meaning be derived without using word order? a comparison of latent semantic analysis and humans. In *Proceedings of Annual Conference of the Cognitive Science Society*, 1997.

Rémi Lebret and Ronan Collobert. Word embeddings through Hellinger PCA. In *Proceedings of EACL*, 2014.

Omer Levy and Yoav Goldberg. Linguistic regularities in sparse and explicit word representations. In *Proceedings of CoNLL*, 2014a.

Omer Levy and Yoav Goldberg. Neural word embedding as implicit matrix factorization. In *Advances in NIPS*, 2014b.

Omer Levy, Yoav Goldberg, and Ido Dagan. Improving distributional similarity with lessons learned from word embeddings. *Trans. ACL*, 3, 2015.

Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. Linguistic regularities in continuous space word representations. In *Proceedings of NAACL-HLT*, 2013a.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *Advances in NIPS*, 2013b.

George A. Miller and Walter G. Charles. Contextual correlates of semantic similarity. *Language and Cognitive Processes*, 6(1), 1991.

Jeff Mitchell and Mirella Lapata. Vector-based models of semantic composition. In *Proceedings of ACL-HLT*, 2008.

Jeff Mitchell and Mirella Lapata. Composition in distributional models of semantics. *Cognitive Science*, 34(8), 2010.

M. A. Montemurro. Beyond the zipf-mandelbrot law in quantitative linguistics. *Physica A*, 300, 2001.

Masayasu Muraoka, Sonse Shimaoka, Kazeto Yamamoto, Yotaro Watanabe, Naoaki Okazaki, and Kentaro Inui. Finding the best model among representative compositional models. In *Proceedings of PACLIC*, 2014.

Partha Niyogi and Federico Girosi. Generalization bounds for function approximation from scattered noisy data. *Advances in Computational Mathematics*, 10, 1999.

Denis Paperno, Nghia The Pham, and Marco Baroni. A practical and linguistically-motivated approach to compositional distributional semantics. In *Proceedings of ACL*, 2014.

Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of EMNLP*, 2014.

J. Pitman. *Combinatorial Stochastic Processes*. Springer-Verlag, 2006.

J. Pitman and M. Yor. The two-parameter poisson-dirichlet distribution derived from a stable subordinator. *Annals of Probability*, 25, 1997.

Sascha Rothe and Hinrich Schütze. Autoextend: Extending word embeddings to embeddings for synsets and lexemes. In *Proceedings of ACL-IJCNLP*, 2015.

Richard Socher, Brody Huval, Christopher D. Manning, and Andrew Y. Ng. Semantic compositionality through recursive matrix-vector spaces. In *Proceedings of EMNLP*, 2012.

Karl Stratos, Michael Collins, and Daniel Hsu. Model-based word embeddings from decompositions of count matrices. In *Proceedings of ACL-IJCNLP*, 2015.

Sho Takase, Naoaki Okazaki, and Kentaro Inui. Composing distributed representations of relational patterns. In *Proceedings of ACL*, 2016.

Yee Whye Teh. A hierarchical bayesian language model based on pitman-yor processes. In *Proceedings of ACL*, 2006.

The BNC Consortium. The british national corpus, version 3 (bnc xml edition). Distributed by Oxford University Computing Services, 2007. URL `http://www.natcorp.ox.ac.uk/`.

Ran Tian, Yusuke Miyao, and Takuya Matsuzaki. Logical inference on dependency-based compositional semantics. In *Proceedings of ACL*, 2014.

Ran Tian, Naoaki Okazaki, and Kentaro Inui. Learning semantically and additively compositional distributional representations. In *Proceedings of ACL*, 2016.

Joseph Turian, Lev-Arie Ratinov, and Yoshua Bengio. Word representations: A simple and general method for semi-supervised learning. In *Proceedings of ACL*, 2010.

Peter D. Turney. Mining the web for synonyms: PMI-IR versus LSA on TOEFL. In *Proceedings of EMCL*, 2001.

Peter D. Turney and Patrick Pantel. From frequency to meaning: Vector space models of semantics. *J. Artif. Int. Res.*, 37(1), 2010.

Vladimir N. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, 1995.

Fabio Massimo Zanzotto, Ioannis Korkontzelos, Francesca Fallucchi, and Suresh Manandhar. Estimating linear models for compositional distributional semantics. In *Proceedings of Coling*, 2010.

George K. Zipf. *The Psychobiology of Language: An Introduction to Dynamic Philology*. M.I.T. Press, 1935.